Annual Conference of ITA
ACITA 2010
MINISTRY OF DEFENCE

The Cognitive Virtues of Dynamic Networks

Trung Dong Huynh, Paul Smart, Dave Braines, Nigel Shadbolt, and Katia Sycara

Overview

In this work, the effect of **dynamic networks** on **distributed problem-solving** was examined using a multi-agent simulation environment. Synthetic agents were tasked with the problem of finding optimal solutions to a specific design problem, and they were allowed to communicate the results of their search efforts to other agents via a dynamically-evolving communication network structure.

The main findings are:

- **Dynamic networks contribute to a better profile of problem-solving performance compared to static networks.**
- **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.

Collaborative Problem Solving

A group of 100 agents were tasked with the **search for optimal design solutions** using a variant of the **NK model** paradigm. An NK model is essentially a

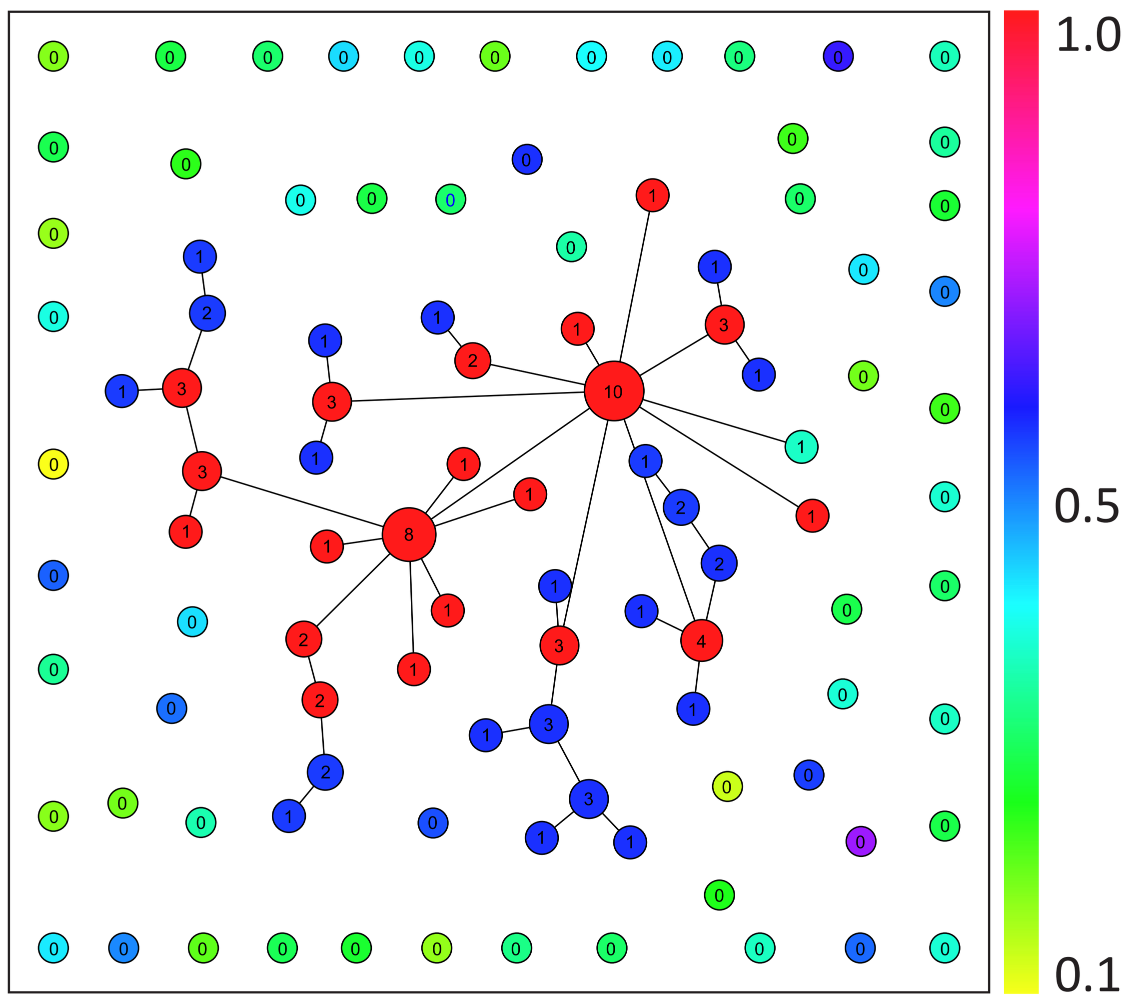

Fig 1: A visualisation of network growth and performance over time. The colours indicate the fitness level of the agents' solutions, with red being the highest fitness score. The above visualisation shows connected agents benefit from superior solutions shared by their connections.

means of generating evolutionary **fitness landscapes** with different degrees of 'ruggedness'. In this task, the agents had to explore a moderately rugged fitness landscape and discover optimal solutions within that landscape.

As part of the task, agents were each assigned **a 20-bit solution string** and their goal was to find more optimal (or fitter) design solutions by progressively modifying the individual elements of the solution string. This is akin to exploring the fitness landscape and **climbing 'up-hill' for optimal solutions**.

We allowed the agents to communicate the results of their search to other agents by organizing them into communication networks with different topologies. If an agent found a better solution from one of its network neighbours, it would **copy the superior solution**.

In our experiments, the agents were connected with each other in

- a **static network** (whose topology was fixed throughout the course of a simulation), or
- a **dynamic network** (where new network links were progressively added to the initially disconnected agents).

The following two factors governed how dynamic networks grew:

- **Network Growth Rate (NGR)**: the rate at which communication links between the agents were added to the network.
- **Network Growth Delay Period (NGDP)**: the period that must elapse before the first link is added to the network.

Variable	Unit	Experimental values						
NGR	link/cycle	0.1	0.2	0.5	1.0	2.0	5.0	10
NGDP	cycle	0	10	20	30	40	50	

A combination of NGR and NGDP results in a two-way factorial design consisting of (7 × 6) 42 experimental conditions. 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).

Dynamic versus Static Networks

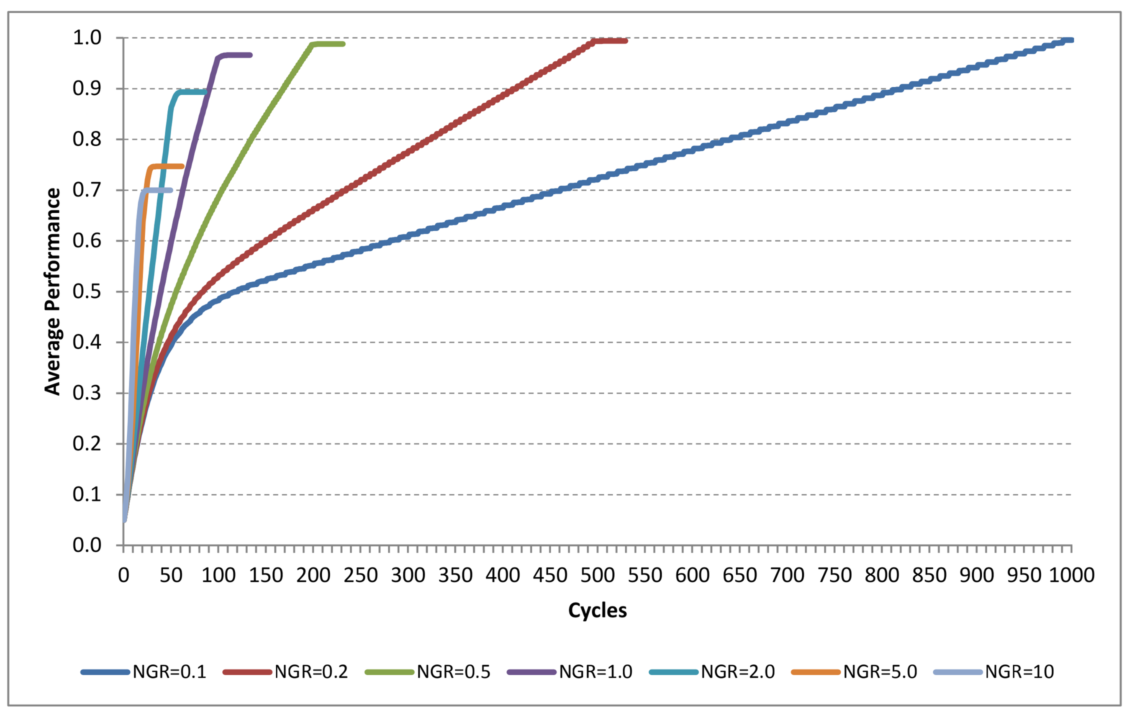

Fig 2: Performance of agents in **dynamic networks** with no initial growth delay period (averaged over 1000 simulations). The faster the network grew, the faster it reached a stable local optimum solution. However, the faster growth rate, at the same time, penalised the quality of the final solution.

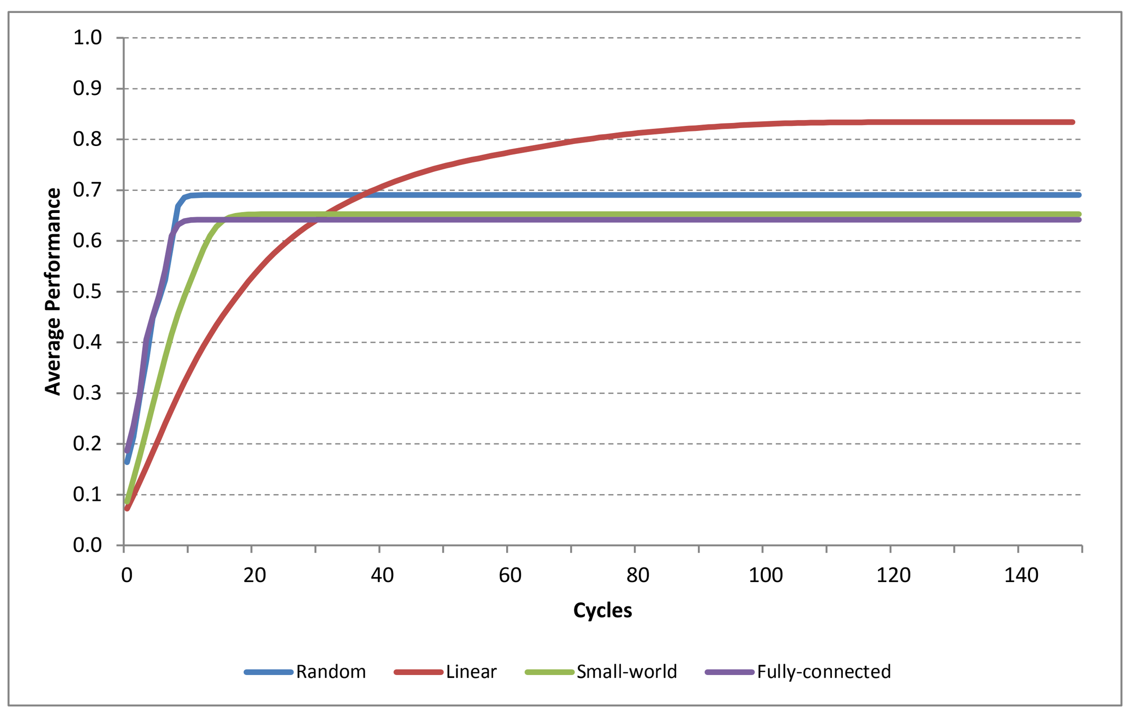

Fig 3: Performance of agents in **static networks** (averaged over 1000 simulations). The static networks settled on a stable local optimum solution much quicker than dynamic networks. However, the final performance in the static networks were significantly lower than that in the dynamic networks (with NGR ≤ 2.0).

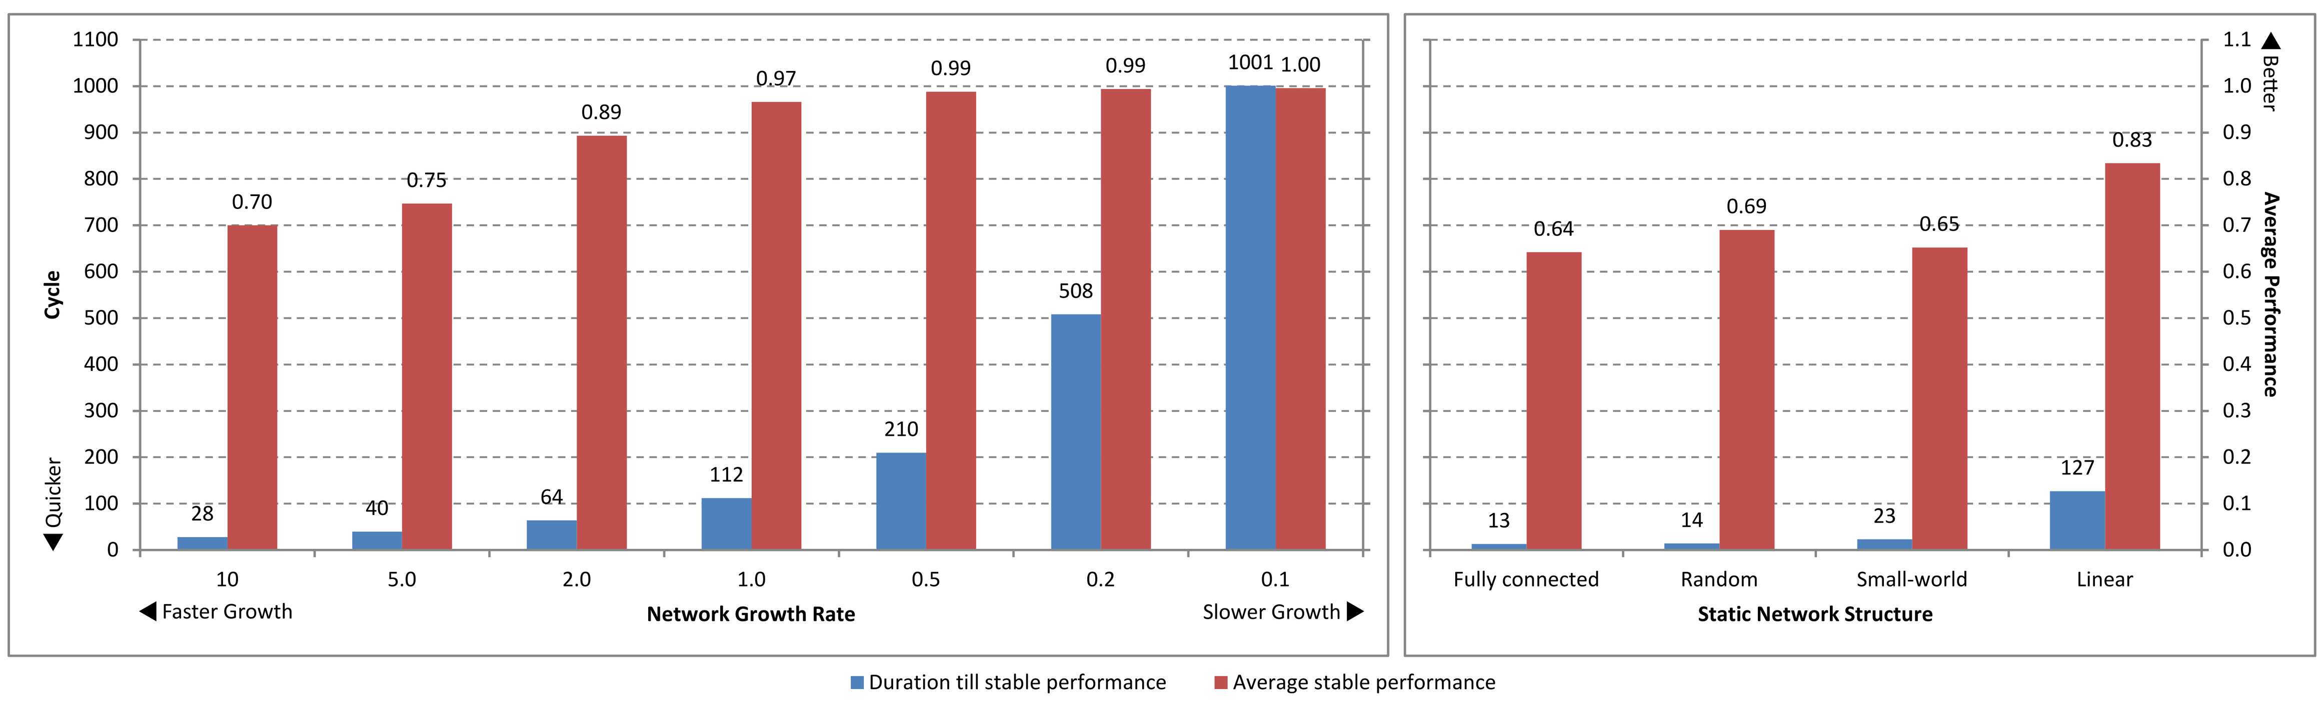

Fig 4: Comparisons of the final performance of networks having different growth rates and static networks (red bars). The adjacent blue bars show the number of cycles it took to achieve the performance indicated by the red bars.

Individual Autonomy versus Social Influence

(Initial Growth Delay Period vs Network Growth Rate)

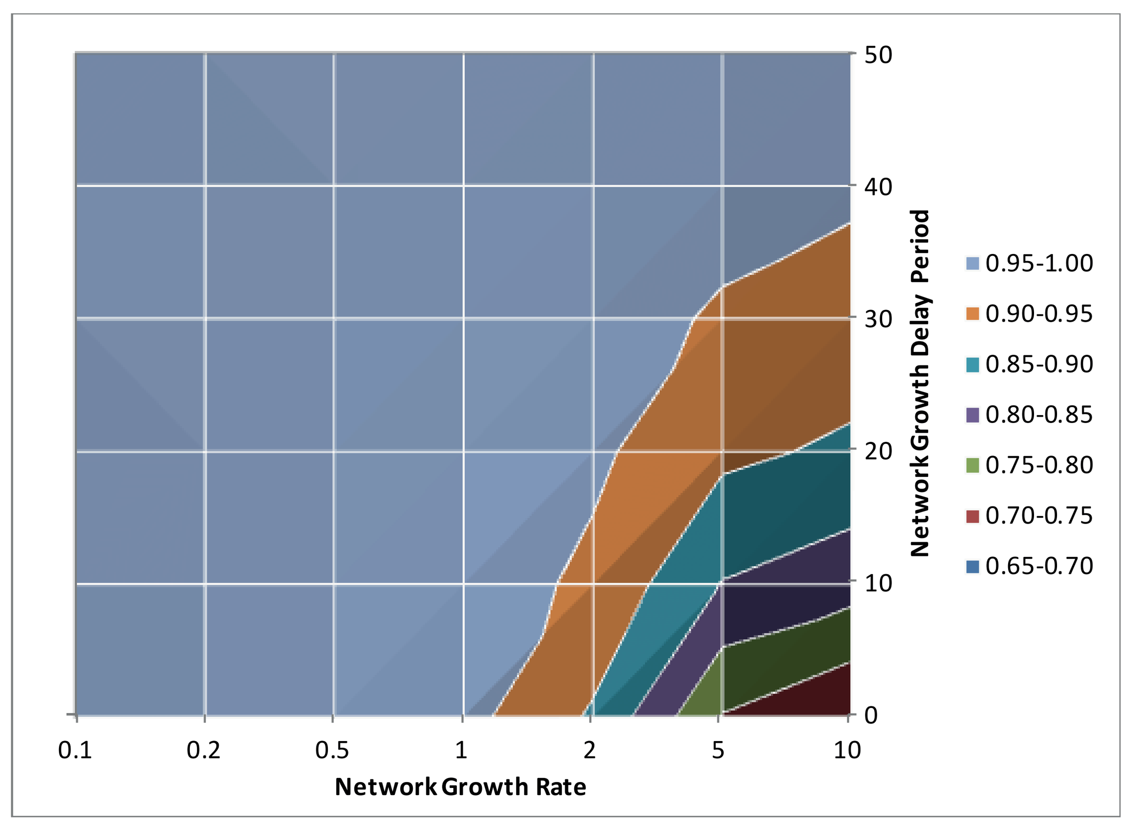

Fig 5: Effects of the various growth delay periods on the final performance of the agents (averaged over 1000 simulations) with different growth rates. Two adjacent colour bands signify a difference of 0.05 point in the average final performance. The growth-rate-related performance penalty was offset by initial delay periods. The faster the growth rate, the longer the initial delay period was needed.