Investigating the Impact of Communication Quality on Evolving Populations of Artificial Life Agents

Sadat Chowdhury[1] and Elizabeth I. Sklar[2]

[1]The Graduate Center, The City University of New York, New York, NY, USA
[2]Department of Computer Science, University of Liverpool, Liverpool, UK
sadatc@gmail.com,e.i.sklar@liverpool.ac.uk

Abstract

The work presented in this paper examines the relationship between *interaction mechanism* and *heterogeneity* in evolutionary multiagent systems. Three interaction mechanisms are explored (stigmergy, broadcast and unicast), in both homogenous and heterogenous populations of artificial agents. A number of different schemes are compared in an experimental environment which concerns agents that evolve to detect, extract and collect resources in a grid-based world. A series of experiments was conducted to analyze the effects of interaction mechanism and population diversity on multiple performance metrics, under various signal transmission qualities. The results show differences in key performance metrics when signal transmission quality degrades, under some experimental conditions. These observations offer an important take-away message for assessing the robustness of systems that may face inconsistent communication network quality, a common problem in today's "wired" real world.

Introduction

In our research, we are investigating the relationship between *interaction mechanism* and *heterogeneity* in evolving multiagent populations. Our work extends earlier experiments conducted by Sklar et al. (2006) in which it was hypothesized that the performance of a multiagent team is affected by both population diversity and level of interaction, such that if interaction mechanisms were employed effectively, a population of heterogenous, single-function agents could generally outperform a homogeneous population of multi-function agents. Several sets of experiments were conducted; and although trends appeared to confirm the hypothesis, the results were statistically inconclusive due to large variances in the experimental data. Here we continue this line of work, in a new version of the simulation environment, *synthScape* (Chowdhury and Sklar, 2015), which runs on a *high-performance computing cluster (HPCC)* so that we can obtain statistically significant results.

The focus of this paper is on evaluating the robustness of interaction mechanisms and population heterogeneity in the face of fluctuation in communication signal. Three interaction mechanisms are explored: *stigmergy (trail)*, *broadcast* and *unicast*. These are tested in two types of evolving populations of artificial agents: *homogenous* and *heterogenous*. The agents collect *resources* from geospatially dispersed locations and bring them to designated *collection sites*. This involves a sequence of *detection*, *extraction* and *transportation* tasks on the part of the agents. The same agent need not be responsible for completing all three tasks; indeed, having the tasks completed simultaneously by different coordinating agents may produce more efficient solutions. Coordination may be achieved by agents *communicating* to indicate that one task in the sequence has been completed at a particular location and the next task can be attempted.

Although the problem space we consider here involves three sequential tasks, our environment and framework are defined to generalise to n-sequential-task domains. The complexity of solutions in sequential-task domains increases with the number of tasks. For example, a foraging problem that requires additional tasks such as *refinement* followed by *processing* of resources, between extraction and transportation, is a more complex 5-task problem than our baseline 3-task version. Each type of task requires a corresponding *trait*, or agent capability, and, in general, solving k-task problems in the sequential-task domain will require k traits. Such problems can be solved by heterogenous populations of simpler specialist agents or homogenous populations of more complex generalist agents. Stated formally, a k-task problem in the sequential task domain can be solved either by: (a) a heterogenous population of n-trait agents where $n \leq k$; or (b) a homogenous population of k-trait agents[1].

The problem of *communication network quality* is real and significant in today's "wifi" society. In the artificial life, multiagent systems and robotics communities, many solutions to various problems are devised and tested in simulation before being deployed an the "real" world. One common problem with transferring solutions to real-world settings is that certain features which work perfectly in simulation are problematic in the real world and can cause the breakdown of an approach that performs well in simula-

[1]When $n > k$, agents have additional capabilities that are not used in solving problems in the domain.

tion. One such feature is *communication*. We have all experienced communication failure when our laptop or mobile phone loses its signal and our connection dies. The work presented here examines the impact of communication quality on several different versions of our evolving multiagent system.

The remainder of this paper is organised as follows. First, we review some related work, specifically research in the literature that has defined taxonomies for interaction or communication mechanisms in multiagent environments. Second, we describe how our experimental environment, *synthScape*, operates. Then we detail our experiments and present results. Finally, we close with discussion.

Related Work

Interactions are an essential requirement for most complex coordination tasks in any *multiagent system (MAS)* (Wooldridge, 2002). Interactions, however, require varying degrees of effort by the agents or specific environmental conditions or both.

Coordination among agents can be achieved using non-interactive techniques, such as tacit agreements, conventions, and simple rules, as has been shown in the famous boid examples in Reynolds (1987). Studies with evolutionary agents, such as Haynes et al. (1995b) and Haynes et al. (1995a), show that simple coordination can be achieved with co-evolved agents without any explicit interactions. The researchers of these studies concluded, however, that such non-interactive techniques are either insufficient or inefficient for multiagent coordination tasks in general. In Campbell et al. (2008), it is shown that non-interactive techniques work well in situations where the ratio of task length to agent team size is small, but their performance decreases as this ratio increases.

More complex coordination among agents can be achieved using indirect forms of interaction mechanisms such as stigmergy. It has been observed that social insects, such as bees and ants, use chemical trails, called *pheromones* that modify the environment and can later be detected by other members of their own species. The modifications (chemical trails) are used as a form of communication; this process of *stigmergy* was first described in Grasse (1959). The concept of stigmergy was used as a basis for agent coordination in techniques such as *ant colony optimization (ACO)* (Dorigo et al., 2000) and *swarm intelligence (SI)* (Dorigo et al., 2006). A real-world use of stigmergy in an industrial application is described in Valckenaers et al. (2004). A physical implementation of stigmergy on real robots was reported in Sahraei et al. (2013).

An even wider range of complex coordination problems can be solved more efficiently with direct interaction mechanisms such as explicit communication. This has been demonstrated in numerous studies (Barlow et al., 2008; Naeini and Ghaziasgar, 2009; Doherty and O'Riordan, 2009). In general, these studies compare the efficiency of overall coordination with and without communication capabilities between the agents and demonstrate the positive aspects of explicit communication.

There have been several attempts to classify interaction mechanisms. In Deugo et al. (2001), several coordination (interaction) mechanisms are presented in the form of software pattern guidelines for developers of MAS. In Menge (1995), and more recently in Eguchi et al. (2006), characteristics that dictate agent interaction behavior is presented; these characteristics can be used as a way to classify interaction mechanisms. Eguchi et al. (2006) lists these characteristics in the form of interaction "attitudes" (shown in Table 1) that an agent (Agent *s*) takes towards another agent (Agent *o*). For example, an agent having a *mutualism* attitude towards another agent will interact with that agent in a way that will benefit both agents.

Attitude	*Agent s*	*Agent o*
Mutualism	Improve	Improve
Harm	Deteriorate	Deteriorate
Predation	Improve	Deteriorate
Altruism	Deteriorate	Improve
Self Improvement	Improve	—
Self Deterioration	Deteriorate	—

Table 1: Interaction Attitudes from Eguchi et al. (2006)

Sklar et al. (2006) lists the following taxonomy classes of interaction mechanisms that form the basis of the interaction mechanisms used in this work:

Tacit agreements: There is no explicit communication between agents, and instead social norms or pre-determined rules govern agent behavior. An example of such an interaction is the flight patterns of flocks of migratory birds: each bird tries to maintain an average distance from other birds and tries to fly in the same general direction as the rest of the flock.

Environmental cues: Agents modify the environment in such a way that another agent can detect and act on that modification. An example of such an interaction is that of an ant leaving a trail of pheromones that others can follow.

Signal broadcasting: Agents explicitly broadcasts signals for other agents to receive. This is basically direct inter-agent communication. Table 2 shows parameter values for these interaction mechanisms.

Additional classifications and taxonomies can be found in Wooldridge (2002) and Weiss (1999). There have been studies where classifications were made based on the role played by the environment (Keil and Goldin, 2006) and the information available in the system (Parunak et al.,

Parameter	Tacit Agreement	Environmental Cues	Signal Broadcasting
Interaction Medium	—	environmental	broadcast
Message Content	—	scalar	binary/ scalar
Message Frequency	—	continual	continual

Table 2: Taxonomy classes of interaction mechanisms in Sklar et al. (2006)

2004). In addition, the evolution of agent signaling from interactions—especially signals with semantic content and language-like properties—has been the subject of a number of studies (Tuci, 2009; Nehaniv, 2000, 2005; Skyrms, 2010).

Durfee (2001, 2004) lists stress factors that need to be addressed when scaling up multiagent systems along various dimensions in order to solve increasingly complex problems (Table 3). *Heterogeneity*, *degree of interaction* and *efficiency*—all dimensions that we consider in our work presented here—are seen as important along the stress factors listed.

Factor	*Dimensions*
Agent Population	quantity, complexity, heterogeneity
Task Environment	degree of interaction, distributivity, dynamics
Solution	efficiency, quality, robustness, overhead limitations

Table 3: Multiagent Coordination Stress Factors from Durfee (2001)

Curran and O'Riordan (2006) examine the effects of interactions in the form of cultural learning on both fitness and diversity. Their results indicate that the addition of such interactions (cultural learning) promotes fitness and significantly increases both genotypic and phenotypic diversity in the population. More recent work on language and communication evolution in embodied agent can be found in (Nolfi and Mirolli, 2009; Szathmáry, 2010).

Here we present a methodical study which investigates the relationship between type of interaction mechanism and population heterogeneity, with respect to communication quality.

Experimental Environment

As mentioned above, our experimental environment is called *synthScape*. This framework was developed in *MASON*, a discrete event simulation platform (Luke et al., 2005). Populations of agents are evolved using *genetic programming (GP)* techniques. The control program (*genotype*) of an individual agent is constructed from a custom extension of the Push language (Spector, 2001; Spector et al., 2005). Push programs are series of instructions that are executed by a stack-based virtual machine. Push was designed specifically for evolutionary computation, and thus has two big advantages: (1) any set of Push instructions makes a valid program, so that genetic reproduction operators cannot produce incomplete or invalid programs; and (2) programs can modify themselves by allowing instructions to manipulate the code stack (where instructions reside), which can potentially introduce complex control strategies. In our implementation, each instruction is atomic and either pushes values onto a code stack, executes instruction(s) from the code stack, or executes domain-specific instructions (i.e., sense, move, communicate). Because of the intense computational requirements necessary to produce statistically significant results, *synthScape* has been designed to be distributed across several processors in a HPCC[2], where each node in the HPCC executes one experimental scenario at a time.

The experiments described in the next section explore communication quality for six different types of populations. Each experiment used the following parameters:

$$\begin{array}{rcrr} \text{dimensions of simulated world} & = & & 16 \times 16 \\ \text{population size} & = & & \{8, 24\} \\ \text{resource capture goal} & = & 12 & (75\%) \\ \text{obstacle density} & = & 32 & (12.5\%) \\ \text{resource density} & = & 16 & (6.25\%) \\ \text{collection site density} & = & 4 & (1.56\%) \end{array}$$

Each population is characterised by: *interaction mechanism* and *heterogeneity*. Populations evolve using an evolutionary algorithm inspired by (Watson et al., 1999; Bianco and Nolfi, 2004), where the concept of evolution is embodied within the agents. An agent maintains its own gene pool, runs its own evolutionary algorithm and evolves its own genotype with a particular set of traits. When the control code (genotype) has finished its execution, it is replaced by the next generation's control code. Agents belonging to the same species, situated in close proximity, are able to copy each others' genotypes to add to their own gene pools. The concept is similar to mating in nature and may allow the transfer of useful genotypes across gene pools and unify the behaviour of agents belonging to the same species. Other scale-related benefits can be realised through an extended form of this basic model (Laredo et al., 2011).

Different types of **interaction mechanisms** facilitate communication between the evolving agents. During an interaction, signals are transferred from agents in "sender" mode to agents in "receiver" mode. The receiving agents

[2] `http://wiki.csi.cuny.edu/cunyhpc/`

may ignore the signals or decide to act upon them. Here, we employ signals with elementary semantics that indicate the state of detected resources and their location. The resources can be in one of three states: FOUND-STATE (thus, ready to be extracted), EXTRACTED-STATE (thus, ready to be processed), or PROCESSED-STATE (thus, ready to be transported). Three interaction mechanisms are considered here: ***trail—***where senders have the ability to leave "trail signals" on the grid, a form of stigmergy; ***broadcast—***where senders have the ability to broadcast signals to receivers within a certain range, who receive the signals instantly; and ***unicast—*** where senders have the ability to send signals to the closest receiver.

Different levels of **heterogeneity** within a population represent the ability of individual agents to perform one or more tasks. We define the heterogeneity of a population based on three factors: *species richness*, *species evenness* and *species diversity*. Species richness, S, is the number of unique species in a population. Species evenness (Peet, 1974; Mulder et al., 2004), E, is based on Simpson's dominance index:

$$E = H/ln(S)$$

where H, the measure of species diversity, is based on Shannon's entropy (Shannon, 1948):

$$H = -\sum_{i=1}^{S} p_i ln(p_i)$$

and p_i is the proportion of species i. E can range from 0, when there is only one dominant species in the population, to 1, when all species are equally abundant in the population. Higher values of H are indicative of both greater richness and evenness in the population; if there is only one species, H approaches 0. The two types of populations considered in the experiments we present here are: ***heterogenous***, where $(S, E, H) = (3, 1, 1.098)$; and ***homogenous***, where $(S, E, H) = (1, 0, 0)$.

Experiments

The experiments conducted here examine the impact of communication quality on six different evolutionary multi-agent populations. The different populations are pairwise combinations of the 3 interaction mechanisms (trail, broadcast and unicast) and 2 population heterogeneity levels (homogeneous and heterogeneous), described in the previous section.

Below are some additional details explaining how each interaction mechanism works:

- *trail:* an agent can drop a trail signal in a cell that any other agent moving through the same cell can detect. Trail signal strength deteriorates over time (currently set to dissipate at the rate of 60% at every step) and does not convey additional semantics to the receiving agent. Receiving agents can: (a) acknowledge that a trail was detected and set an internal state variable or (b) move towards a neighboring cell containing the highest concentration of signal strength.
- *broadcast:* an agent can send one of two signals: SIGNAL-A and SIGNAL-B. Once a signal is transmitted, any agent can receive that signal within 2 time steps of the simulation. Receiving agents can: (a) acknowledge it and set an internal state variable or (b) move towards the source of the signal.
- *unicast:* an agent can send one of two signals: SIGNAL-A and SIGNAL-B to the agent that is closest to it. Receiving agents can: (a) acknowledge it and set an internal state variable or (b) move towards the source of the signal.

Note that the meaning of the generic SIGNAL-A and SIGNAL-B depends on which agent transmits the signal (see below).

A number of *constraints* were placed on agents' actions when the broadcast and unicast interaction were used. These are as follows:

- An agent with detection capability is only allowed to send SIGNAL-A.
- An agent with the extraction capability is only allowed to extract a resource once it has received SIGNAL-A. An extractor agent may also send SIGNAL-B to announce that there is an extracted resource that is ready to be collected.
- An agent with the transportation capability is only allowed to load, carry and unload an extracted resource—potentially to a collection site—once it has received SIGNAL-B.

These constraints ensure that the agents' actions are instigated as the result of receiving a communication signal, as opposed to occurring randomly (e.g., because an extractor agent randomly stumbled on a cell where a found resource was ready to be extracted).

Within each of the six interaction-mechanism/population-heterogeneity combinations evaluated, the communication signal quality was varied. Two extreme values were tested: 100% was the highest quality, meaning that no signals were lost; and 25% was the poorest quality, meaning that three-quarters of signals transmitted were lost. Two intermediate values, 75% and 50%, were also tested, indicating one-quarter and one-half of the signals transmitted were lost, respectively.

Results

The following section describes the results of our experiments. Each of the 6 populations tested was evolved for 1000 generations. Using the HPCC, 100 runs were executed for each. Average results are presented.

We examine the results in terms of 4 metrics: (a) resource capture rate; (b) interaction instruction transmission rate; (c) signals received; and (d) resource collection interval. Each is presented below.

Resource Capture Rate. Figure 1 contains the *resource capture rates* for all 6 experimental conditions. This metric indicates the percentage of resources in the agents' environment which were successfully detected, extracted and transported. We make several observations about our results. First, both homogenous and heterogenous populations evolve to near 100% capture rate when interacting with trail signals. Second, the heterogenous agents evolve towards optimal capture rate much faster than the homogenous agents. This is more apparent when the agents communicate using broadcast and unicast. Third, the behavior of both populations using the *trail* interaction mechanism are very similar. One subtle difference is that the heterogenous agents achieve 50% capture rate within 150 generations, whereas it takes close to 300 generations in the case of the homogenous agents. Both populations seem to be not very sensitive to the signal transmission quality. Fourth, in both populations, signal transmission quality in the *broadcast* interaction mechanism positively impacts the capture rate. Overall, the heterogenous agents perform better than the homogenous agents (seemingly by some constant factor). Fifth, the performance of both populations with the *unicast* interaction mechanism mirrors that of broadcast, but seems to be reduced by some constant factor in both populations.

Next, we examine the ranking of the results. The best performance came from heterogenous agents using trails, irrespective of the quality of the communication signal. The second-best performance came from homogenous agents using trails, again irrespective of the quality of signal. The third-best performance came from heterogenous agents using broadcast with 100% signal transmission quality. The fourth-place performance came from heterogenous agents using unicast with 100% signal transmission quality. The worst performance came from homogenous agents using broadcast and unicast with 25% signal transmission quality.

In conclusion, with respect to resource capture rate, the heterogenous agents with trail perform best overall. There might be situations where trail is not an ideal form of communication. For example, trails can be arbitrarily removed by the environment, and there can be "cross-talk" (where different signals conflict).

Interaction Instruction Transmission Rate. Figure 2 plots the number of interaction instructions—send and receive commands—that were transmitted by the agents as they were evolving. We make several observations about our results. First, the heterogenous agents transmit significantly higher numbers of interaction instructions. In the case of trail, homogenous agents evolve to practically not transmitting any send or receive trails—even though they can utilize the signals. Second, the heterogenous-trail agents settle on a fairly consistent transmission rate (e.g., 100 instructions at the signal transmission rate of 100%). Third, in both populations using the broadcast interaction mechanism, signal transmission quality positively impacts the rate of issuing interaction. Overall, the heterogenous agents issue significantly more instructions than homogenous agents. Fourth, with the heterogenous-unicast agents, the signal transmission quality seems to have little or no impact on the rate of transmissing interaction instructions. Although the signal transmission quality is somewhat differentiated by the homogenous population, it issues significantly fewer instructions. Fifth, the number of instructions transmitted seems to steadily rise over generations for both broadcast and unicast interaction mechanisms. Sixth, in the heterogenous populations, there is an initial spike, followed by a sharp decline, and then a steadier rise of interaction instructions. One possible explanation for this behavior is that in the initial generations, communication behavior provides a high fitness and spreads among the population; however, there is a sharp decline once the population becomes more selective in its use of communication instructions. In the homogenous population, where communication plays less role, this dipping behavior is absent.

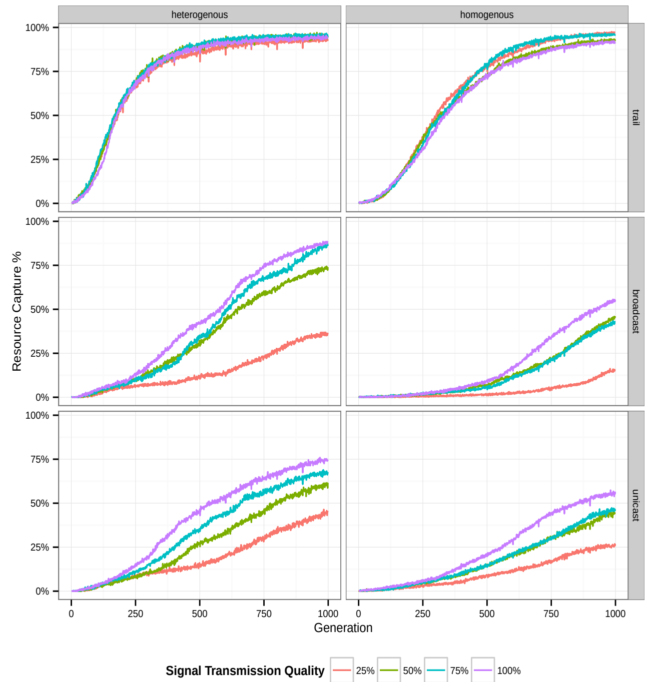

Figure 1: Resource Capture rate (y-axis).

Next, we examine the ranking of the results. The best results were obtained by heterogenous agents using broadcast with higher signal transmission rates. The second-best results were obtained by heterogenous agents using unicast, irrespective of the signal transmission rate.

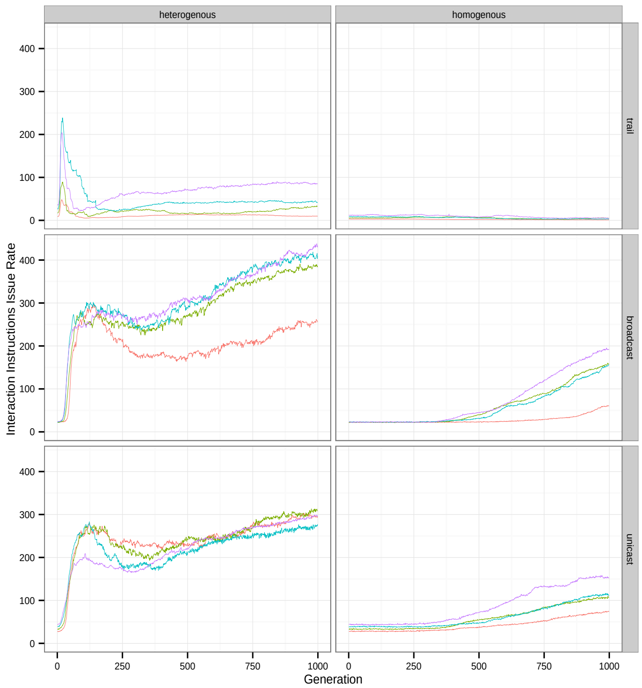

Figure 2: Transmission Rate (y-axis). Legend is the same as in Figure 1.

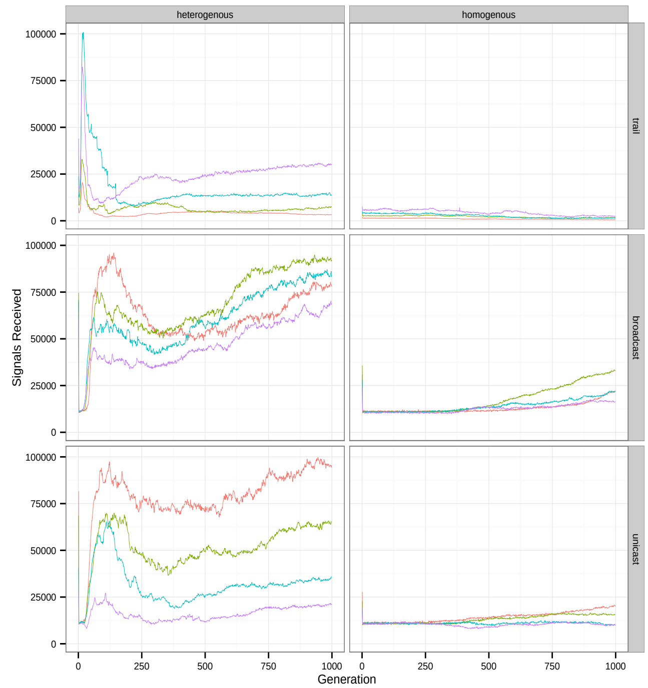

Figure 3: Signals Received (y-axis). Legend is the same as in Figure 1.

Signals Received. Figure 3 contains the number of signals that are actually received in each generation. We make several observations about our results. First, for all interaction mechanisms, heterogenous agents evolve to receive more signals than the homogeneous agents. Second, when heterogenous agents use the trail mechanism, they evolve to receive more signals when the signal transmission rate is higher; whereas, in the case of broadcast and unicast, the agents evolve to receive fewer signals. This suggests that with trail, the receipt of more signals results in better optimization, whereas with broadcast and unicast, receipt of fewer signals results in better optimization.

Resource Collection Interval. Figure 4 shows the evolution of the average number of steps between resource collections. We make several observations about our results. First, the number of steps settles down to a smooth curve very quickly in the case of heterogenous agents for all interaction mechanisms, as compared to homogeneous agents. Second, the plots show how much the transmission quality impacts the average collection interval in both homogenous and heterogenous agents. In all cases, a high transmission quality significantly lowers the collection interval and keeps the value low across the generation (there is less jitter). Third, the heterogenous-trail agents settle down to a consistently lower interval much more quickly, regardless of the transmission quality.

Next we examine the ranking of the results. The best results were obtained by the heterogenous and homogenous agents using the trail interaction mechanism. The second-best results were obtained by the heterogenous-unicast agents with 100% transmission quality. The third-best results were obtained by the heterogenous-broadcast agents with 100% transmission quality.

Discussion

A brief discussion of our overall results follows.

First, the best task completion rates (capture rate) are evolved by both the homogenous and heterogenous agents using the trail interaction mechanism. Additionally, heterogenous agents evolve to higher efficiency levels faster than homogenous agents.

Second, homogenous agents evolve to use much less communication than heterogeneous agents. Our conclusion is that communication adds extra overhead that homogenous agents tend to keep to a minimum, because they can perform all the tasks themselves (i.e., without needing the cooperation of other agents, like the heterogenous agents do).

Third, in this constrained version of our environment, heterogenous agents are forced to communicate in order to capture resources. The result is that they achieve higher efficiency levels than in our previous work, which was conducted in an unconstrained version of our environment (Chowdhury and Sklar, 2015).

Fourth, the signal transmission quality overall has a pos-

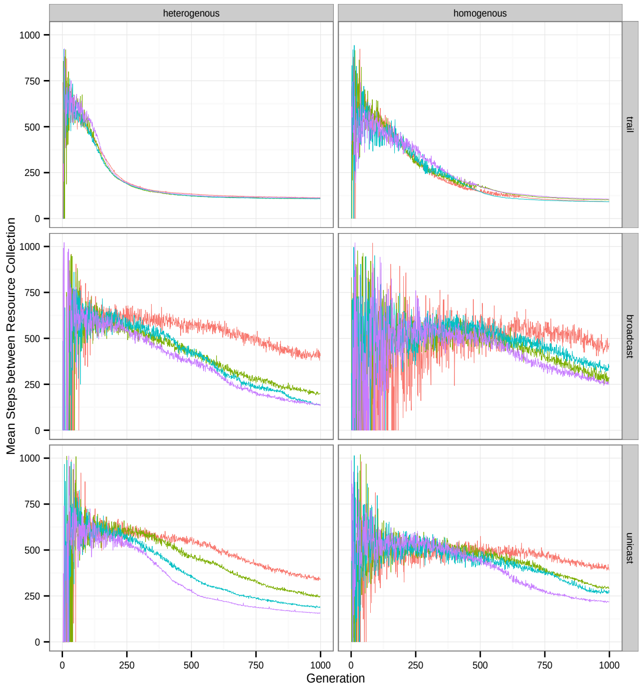

Figure 4: Interval Between Resource Collections. Legend is the same as in Figure 1.

itive impact on the task completion rate. In other words, the better the signal transmission quality, the higher the task completion rate.

Fifth, broadcast and unicast receivers evolve to spend less instructions to receive signals when the transmission rates are better. In the case of the trail mechanism, when transmission rates are better, more "receive signal" instructions are issued. This makes sense: trail signals do not convey any semantics—they are generic messages. In contrast, broadcast and unicast signals convey richer semantics: for example, whether something is an extractable resource.

Finally, aside from the trail mechanism, unicast with 100% signal transmission quality evolves the best (lowest) resource collection interval. This also makes sense: unicast is sent to the closest agent and this should theoretically lower the collection times.

In conclusion, we have shown that the communication quality has an impact on key performance metrics in an evolutionary multiagent environment applied to a 3-task sequential domain. However, the overall impact of communication is not necessarily positive and is dependent on the heterogeneity of the population. The impact is more significant with broadcast and unicast modes of interaction in the heterogenous populations; these interaction mechanisms are richer from a semantic standpoint with respect to message content. The trail mode of interaction is more resilient to degradation in quality of communication in both heterogenous and homogenous populations.

Our next steps with this line of work will examine how diversity impacts performance, for example how various species diversity, evenness and richness values effect performance.

Acknowledgements

This research was supported, in part, by a grant of computer time from the City University of New York High Performance Computing Center under NSF Grants CNS-0855217, CNS-0958379 and ACI-1126113; as well as a University of Liverpool Research Fellowship.

References

Barlow, G. J., Oh, C. K., and Smith, S. F. (2008). Evolving cooperative control on sparsely distributed tasks for UAV teams without global communication. In Keijzer, M., Antoniol, G., Congdon, C. B., Deb, K., Doerr, B., Hansen, N., Holmes, J. H., Hornby, G. S., Howard, D., Kennedy, J., Kumar, S., Lobo, F. G., Miller, J. F., Moore, J., Neumann, F., Pelikan, M., Pollack, J., Sastry, K., Stanley, K., Stoica, A., Talbi, E.-G., and Wegener, I., editors, *GECCO '08: Proceedings of the 10th annual conference on Genetic and evolutionary computation*, pages 177–184, Atlanta, GA, USA. ACM.

Bianco, R. and Nolfi, S. (2004). Toward open-ended evolutionary robotics: evolving elementary robotic units able to self-assemble and self-reproduce. *Connect. Sci.*, 16(4):227–248.

Campbell, A., Wu, A. S., and Shumaker, R. (2008). Multi-agent task allocation: Learning when to say no. In *Proceedings of the 10th Annual Conference on Genetic and Evolutionary Computation*, GECCO '08, pages 201–208, New York, NY, USA. ACM.

Chowdhury, S. and Sklar, E. I. (2015). Exploring interaction, diversity and efficiency of biologically inspired evolutionary multiagent systems. In *8th International Conference on Bio-inspired Information and Communications Technologies (formerly BIONETICS)*. ICST.

Curran, D. and O'Riordan, C. (2006). Increasing population diversity through cultural learning. *Adaptive Behavior - Animals, Animats, Software Agents, Robots, Adaptive Systems*, 14(4):315–338.

Deugo, D., Weiss, M., Kendall, E., et al. (2001). Reusable patterns for agent coordination. *Coordination of Internet Agents, Springer*.

Doherty, D. and O'Riordan, C. (2009). Effects of shared perception on the evolution of squad behaviors. *IEEE Transactions on Computational Intelligence and AI in Games*, 1(1):50–62.

Dorigo, M., Birattari, M., and Stutzle, T. (2006). Ant colony optimization. *Computational Intelligence Magazine, IEEE*, 1(4):28–39.

Dorigo, M., Bonabeau, E., and Theraulaz, G. (2000). Ant algorithms and stigmergy. *Future Generation Computer Systems*, 16(8):851–871.

Durfee, E. (2001). Scaling up agent coordination strategies. *Computer*, 34(7):39–46.

Durfee, E. (2004). Challenges to scaling-up agent coordination strategies. *An Application Science for Multi-Agent Systems*, pages 113–132.

Eguchi, T., Hirasawa, K., Hu, J., and Ota, N. (2006). A study of evolutionary multiagent models based on symbiosis. *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, 36(1):179–193.

Grasse, P. P. (1959). La reconstruction du nid et les coordinations interindividuelles chez bellicositermes natalensis et cubitermes sp. la theorie de la stigmergie: essai d'interpretation du comportement des termites constructeurs. *Insectes Sociaux*, 6:41–81.

Haynes, T., Sen, S., Schoenefeld, D., and Wainwright, R. (1995a). Evolving multiagent coordination strategies with genetic programming. Technical Report UTULSA-MCS-95-04, The University of Tulsa.

Haynes, T., Sen, S., Sen, I., Schoenefeld, D., and Wainwright, R. (1995b). Evolving a team. In *In Working Notes of the AAAI-95 Fall Symposium on Genetic Programming*, pages 23–30. AAAI.

Keil, D. and Goldin, D. (2006). Indirect interaction in environments for multi-agent systems. In *Proceedings of the 2nd international conference on Environments for Multi-Agent Systems*, E4MAS'05, pages 68–87, Berlin, Heidelberg. Springer-Verlag.

Laredo, J. L. J., Lombraña González, D., Fernández de Vega, F., Arenas, M. G., and Merelo Guervós, J. J. (2011). A peer-to-peer approach to genetic programming. In Silva, S., Foster, J. A., Nicolau, M., Giacobini, M., and Machado, P., editors, *Proceedings of the 14th European Conference on Genetic Programming, EuroGP 2011*, volume 6621 of *LNCS*, pages 109–118, Turin, Italy. Springer Verlag.

Luke, S., Cioffi-Revilla, C., Panait, L., Sullivan, K., and Balan, G. (2005). Mason: A multiagent simulation environment. *Simulation*, 81(7):517–527.

Menge, B. (1995). Indirect effects in marine rocky intertidal interaction webs: patterns and importance. *Ecological monographs*, 65(1):21–74.

Mulder, C. P. H., Bazeley-White, E., Dimitrakopoulos, P. G., Hector, A., Scherer-Lorenzen, M., and Schmid, B. (2004). Species evenness and productivity in experimental plant communities. *Oikos*, 107(1):50–63.

Naeini, A. T. and Ghaziasgar, M. (2009). Improving coordination via emergent communication in cooperative multiagent systems: A genetic network programming approach. In *IEEE International Conference on Systems, Man and Cybernetics, 2009. SMC 2009*, pages 589–594.

Nehaniv, C. (2000). The making of meaning in societies: Semiotic and information-theoretic background to the evolution of communication. *Society for the Study of Artificial Intelligence and Adaptive Behavior*, pages 73–84.

Nehaniv, C. L. (2005). Open problems in the emergence and evolution of linguistic communication: A road-map for research. In *Proc. Second International Symposium on the Emergence and Evolution of Linguistic Communication*, pages 86–93.

Nolfi, S. and Mirolli, M. (2009). *Evolution of communication and language in embodied agents*. Springer Science & Business Media.

Parunak, H. V. D., Brueckner, S., Fleischer, M., and Odell, J. (2004). A design taxonomy of multi-agent interactions. *Agent-Oriented Software Engineering IV*, pages 123–137.

Peet, R. K. (1974). The measurement of species diversity. *Annual review of ecology and systematics*, 5:285–307.

Reynolds, C. W. (1987). Flocks, herds, and schools: A distributed behavioral model. *SIGGRAPH Computer Graphics*, 21(4):25–34.

Sahraei, B. R., Alers, S., Tuyls, K., and Weiss, G. (2013). Stico in action. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1403–1404, Saint Paul, MN, USA.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell system technical journal*, 27.

Sklar, E., Schut, M., Diwold, K., and Parsons, S. (2006). Exploring coordination properties within populations of distributed agents. *AAAI 2006 Spring Symposium on Distributed Plan and Schedule Management*.

Skyrms, B. (2010). *Signals: Evolution, learning, and information*. OUP Oxford.

Spector, L. (2001). Autoconstructive evolution: Push, pushgp, and pushpop. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2001)*, pages 137–146.

Spector, L., Klein, J., and Keijzer, M. (2005). The push3 execution stack and the evolution of control. In *Proceedings of the 2005 conference on Genetic and evolutionary computation*, pages 1689–1696. ACM.

Szathmáry, E. (2010). Evolution of language as one of the major evolutionary transitions. In *Evolution of communication and language in embodied agents*, pages 37–53. Springer.

Tuci, E. (2009). An investigation of the evolutionary origin of reciprocal communication using simulated autonomous agents. *Biol. Cybern.*, 101(3):183–199.

Valckenaers, P., Kollingbaum, M., Van Brussel, H., et al. (2004). Multi-agent coordination and control using stigmergy. *Computers in Industry*, 53(1):75–96.

Watson, R. A., Ficici, S. G., and Pollack, J. B. (1999). Embodied evolution: Embodying an evolutionary algorithm in a population of robots. In *Proceedings of the Congress on Evolutionary Computation*, volume 1, pages 335–342, Mayflower Hotel, Washington D.C., USA. IEEE Press.

Weiss, G., editor (1999). *Multiagent systems: a modern approach to distributed artificial intelligence*. MIT Press, Cambridge, MA, USA.

Wooldridge, M. (2002). *An Introduction to Multi-Agent Systems*. John Wiley & Sons, Inc., New York, NY, USA.