

Comparing Communication Protocols under Cooperative Pressure

A.E. Eiben M.C. Schut N. Vink

Artificial Intelligence Section
Department of Computer Science
Vrije Universiteit Amsterdam (NL)
{gusz,schut}@cs.vu.nl nvink@illyan.nl

Abstract. The effectiveness of coordination within a multi-agent system depends largely on the abilities of the agents to cooperate and communicate. In environments where coordination is essential for goal accomplishment, agents must thus appropriately communicate for cooperative task achievement. We investigate two communication protocols (centralised and distributed) in learning and non-learning agent societies. This work is part of a larger project whose main goal is to investigate the emergence of cooperation and communication in response of (scalable) environmental challenges.

1 Introduction

For evolutionary simulation, it is recognised that the simulated environment should be a first order object as it guides the selection and adaptation process [5]. The functionality of such an environment for the agents is defined as a shared domain through which they communicate, (inter)act or both [7]. The motivation for our work presented here is the so-called *information speed hypothesis* [2]. It says that there is a strong dependency between the speeds of two agent-environment interaction processes: 1) the outdateding of information that agents act upon, and 2) the communication by which agents exchange information. For example, a gossip may spread quickly, but may become obsolete even faster. We are particularly interested in properties of the environment (determining the speed at which information becomes obsolete) in relation to properties of the agent communication graph (e.g., shortest path, average out degree) with respect to the hypothesis.

The main research objective of this paper is to compare the effects of two drastically different communication protocols in a wide variety of artificial agent societies using these protocols. Firstly, we investigate a *centralised* communication method, where individuals multicast messages that can be received by any individual. Technically, this method was implemented through a message board where senders can put their messages and any individual can read the present content of this message board. Secondly, we investigate a *decentralised* communication method where information is transferred directly between agents without a third party (messageboard, or alike). In comparison with [2], we considerably extend the work presented there as we investigate a much wider range of environments in this paper.

In a series of experiments, we create a range of agent environments by varying the available *resources* and the *cooperation threshold*; the resources are needed by the agents to survive, the cooperation threshold determines how strong the pressure is on the agents to cooperate. Our agents are equipped with a hardwired mechanism for communication, and learn (by evolution) to use this mechanism.

2 Background

Messageboard Communication In the centralised approach, communication between agents is supported by a centralised component; this is a component that is accessible by all agents, e.g., communication through a *messageboard*. A messageboard enables the agents to communicate by facilitating the storage and retrieval of communicated information. All agents wanting to communicate can access the board to post a message to it, or read messages posted on it.

Newscast Communication In the decentralised approach, there is no central support for communication. Agents that wish to communicate to other agents need to ‘manually’ find agents to communicate to and exchange the information with them. Many decentralised communication protocols (e.g., gossip, epidemic-based) have been proposed and researched recently. We have implemented the *newscast* model in the experiments described below [4]. The essential part of the protocol is the agent *cache*. Each agent holds in its cache a collection of messages it has received from other agents. When an agent wants to spread information, it 1) randomly picks a agent ID from its cache, 2) adds the information it wants to communicate to its own cache, 3) timestamps and puts its ID on the message, and 4) sends the information. The caches of the communicating agents are merged, after which so many old messages are removed such that it fits in the cache. Note that the caches of the agents define the earlier mentioned communication graph. Simulations from [4] show that the newscast model is a robust and scalable solution for problems in which decentralised information dissemination is needed.

3 System Description

The system in which we conduct our experiments consists of a simulated environment that represents an artificial society (called VUSCAPE) and a set of agents that populates this society.

Simulated Environment The simulated environment is an artificial society called VUSCAPE [1], which is based on SUGARSCAPE [3]. This artificial society concerns a two dimensional grid, wrapped around the edges, where each position corresponds with an area which can contain multiple agents and some amount of sugar. Agents move through the world by vertically or horizontally jumping to another location (cf. moving in SUGARSCAPE). The agents live off the sugar, determining their level of fitness; if an agent’s fitness reaches zero, it dies. The major differences between VUScape and SUGARSCAPE concern the implementations of cooperation, communication, explorative behaviour, increased grid-point inhabitation, randomised sugar distribution, and




Fig. 1. The agent control loop.

randomised age initialisation. We investigated the effects of these differences experimentally in [1].

Environment - Cooperation Each agent can only harvest a maximum amount of sugar on its own. This amount is set by the `maxSugarHarvest` (MSH) parameter. If an agent is at a location at which the amount of sugar is over this threshold, it needs other agents to harvest the sugar. If there are more agents at such a location, these agents harvest the sugar together and the sugar is evenly distributed over these agents. In the experiments described below, the cooperation threshold is the same for all agents. In addition to the MSH parameter, we scale the necessity to cooperate by varying the number of available resources in the environment, called the maximum sugar size (MSS) in VUSCAPE. Based on the settings of MSH (implements the earlier mentioned *cooperation threshold*) and MSS (implements the earlier mentioned *available resources*), we can create easier and more difficult environments for the agents to survive.

Agents Our agents were based on the classical SUGARSCAPE agent design: prominent features include metabolism, gender, child bearing, death, vision, allow sex and replacement. An agent was able to detect agents and resources for a number of grid-cells determined by its *vision* parameter. It was able to move for a number of grid-cells determined by its *move* parameter. The control loop of the agent is shown in Figure 1. This loop is explained in detail in [1].

Agents - Communication Agents are endowed with talk and listen capabilities. The talk feature determines whether the agent performs a communicative action itself, namely informing other agents of: 1) the amount of sugar that is on its location, and 2) the coordinates of its location. The listen feature is used in the observation and decision making processes of the agent. By listening, the agent receives information from other agents about amounts of sugar at the locations of those agents.

Agents - Evolution Agents underwent an evolutionary process of selection and variation. Agents with a high fitness were selected for and variation of agents was accomplished by *crossover* of agent genotypes. The agent genes involved are the talk and listen preferences and the initial amount of sugar. Crossover happens by reproduction of two agents; this is not subject to individual decisions, nor is there any mate selection. If 1) two agents are on the same location or next to each other, 2) the genders differ, 3) their sugar levels are above the reproduction threshold, 4) they are both in the fertile age range, 5) there is a vacant cell in the vicinity for placing the offspring, agents will always mate and generate offspring. Reproduction takes place by crossover applied to two parents yielding the child, followed by a mutation operator on the child. The talk

Experiment		Agent	
numberOfRuns	10	maxSugarHarvest (msh)	exp
Scape		singleStep	false
height	50	initAgeZero	false
width	50	minVision	1
runLength	2000	maxVision	1
reseedSugar	true	minSugarMetabolism	1
initialPopulation	400	maxSugarMetabolism	1
sugarSeed.uniqueCell	false	minDeathAge	60
sugarDistributionUnif	3	maxDeathAge	100
sugarGrowBackRate	1.0	sexRecoveryPeriod	0
numberOfSeeds	1000	minReproductionSugar	0
sugarDistributionType	uniform	allowSex	true
maxSugarSize (mss)	exp	minInitialSugar	50
Cell		maxInitialSugar	100
allowMultipleAgents	true	preferNearestCell	false

Table 1. An overview of the experimental parameters. Parameters indicated with **exp** are varied in the experiments.

and listen genes express probabilities and are formally real valued numbers between 0 and 1. The value of a gene in a child is the inherited value (from the wealthiest parent) plus a random number drawn from a Gaussian distribution with zero mean and fixed standard deviation σ . The child receives from each of the parents half of their sugar. The child inherits each of the values for vision, age of death, metabolism, and child bearing independently from one of the parents without change. After mating, each agent has a so-called recovery period, which is the number of cycles after mating that an agent cannot mate.

4 Experiments

Setup We conducted two series of experiments, each consisting of 10 independent runs: one experiment with agents that always use their communication capabilities, and another series with agents that evolutionary learn using their communication capabilities. For both series, we conducted a number of runs with centralised (messageboard) communication and a number with decentralised (newscast) communication. An overview of all experimental parameter values is given in Table 1. Additionally, talk and listen features are inherited from the parent with the most sugar. The mutation sigma is 0.1.

Results The obtained results have been included in Figure 2. It shows the comparison of protocols where the average population sizes are measured against the number of available resources (maximum sugarsize - MSS) and the cooperation threshold (maximum sugar harvest - MSH). Each point in these is the average outcome (average population size over 2000 iterations) of 10 independent runs.

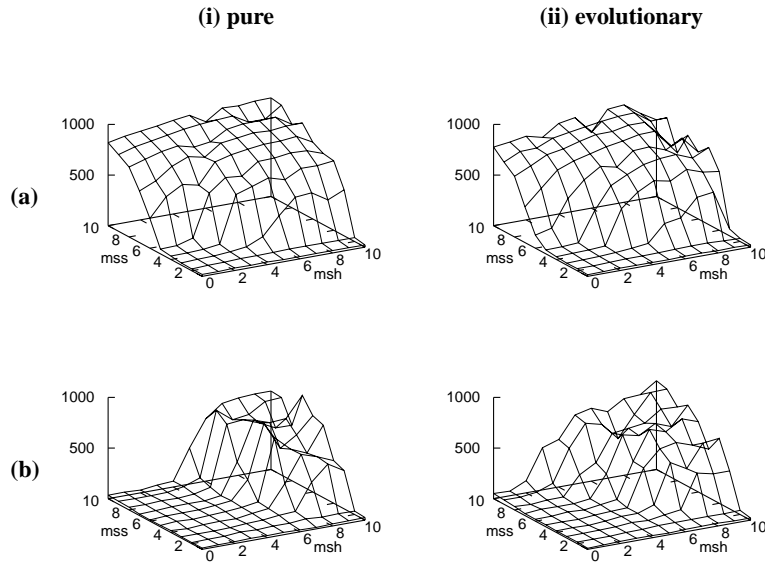


Fig. 2. Experimental results, where the average population sizes are measured against the number of available resources (maximum sugarsize mss) and the cooperation threshold (maximum sugar harvest). Row (a) shows the results with the messageboard protocol and row (b) with the newscast protocol.

5 Findings

(All references are to Figure 2.) For the pure variants (comparing a.i and b.i), we see that the messageboard protocol (a) performs much better than the newscast protocol (b). Also for the evolutionary variant (compare a.ii and b.ii), we see that the messageboard protocol (a) outperforms the newscast protocol (b). Note however that the difference is smaller, because the performance of the newscast protocol has improved (compare b.i and b.ii). In general, comparing i and ii, we see that the messageboard performs better in the pure variant than in the evolutionary variant (compare a.i and a.ii), whereas the newscast protocol performs better in the evolutionary variant (compare b.i and b.ii).

We conducted an additional experiment (of which the results have not been included here because of space restrictions) without communication (more details, including results, in [6]). The pure newscast protocol (b.i) gives little advantage over no communication, the evolutionary newscast (b.ii) gives considerable advantage, and the messageboard protocol (both a.i and a.ii) gives much advantage.

An earlier considered explanation of related experimental results concerning this topic [2] is confirmed here. In the VUSCAPE environment, communicated information outdates extremely fast, i.e., the probability that sugar at some location heard of has not been yet consumed when the listener arrives there, is very low. This means that the

communication protocol must be able to distribute and update the information at such rate that still the agents receive accurate information. We observe that the messageboard can do this whereas the newscast cannot. It is very plausible that our results can thus be explained by analysis of a property of the environment rather than the used protocols.

Overall, we can say that in the VUSCAPE environment 1) communication is advantageous to agents who use it, although 2) if agents can survive without it, the added value of communication is relatively small.

6 Conclusions

We compared the messageboard protocol with the newscast protocol in a simulated environment representing an artificial society. The results show that the performance of the messageboard protocol is much better than that of the newscast protocol in the environments that we examined. We observe this when the agents always communicate as well as when they have to learn to communicate (by evolution).

Concerning the information speed hypothesis [2], we conclude that the VUSCAPE environment is one in which information outdates very rapidly. The messageboard protocol can deal with this (messages may be directly removed after listened to) whereas the newscast protocol cannot (because of its decentralised nature). Further experimentation may shed more light on this issue, but we expect the (de)centralised nature of agent communication critical with respect to which environments can be effectively dealt with.

References

- [1] Buzing, P.C., Eiben, A.E., and Schut, M.C. *Evolving Agent Societies with VUScape*. In: Banzhaf, W., Christaller, T., Dittrich, P., Kim, J.T., and Ziegler, J., editors. *Advances in Artificial Life, Proceedings of the 7th European Conference on Artificial Life (ECAL 2003)*, Lecture Notes in Artificial Intelligence, volume 2801, Springer Verlag, 2003, pp. 434-441.
- [2] Eiben, A.E., Schut, M.C. and Toma, T., *Comparing Multicast and Newscast Communication in Evolving Agent Societies*. In: *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2005)*, 2005. In press.
- [3] Epstein, J., and Axtell, R. *Growing Artificial Societies*, Brookings Institute Press, 1996.
- [4] Jelasity, M., and Van Steen, M. *Large-Scale Newscast Computing on the Internet*. Technical Report IR-503, Department of Computer Science, Vrije Universiteit Amsterdam, The Netherlands, 2002.
- [5] Klugl, F., Fehler, M., and Herrler, R., *About the Role of the Environment in Multi-agent Simulations*. In: Weyns, D. et al., editors. *Proceedings of the first Workshop on Environments for Multi-Agent Systems (E4MAS 2004)*, Lecture Notes in Artificial Intelligence, volume 3374, Springer Verlag, 2004, pp. 127-149.
- [6] Vink, N. *Exploring Communication Dynamics in VUSCAPE*, Technical report, Department of Computer Science, Vrije Universiteit Amsterdam, The Netherlands, 2004. Available online at <http://www.cs.vu.nl/ci/eci/>.
- [7] Weyns, D., Van Dyke Parunak, H., Michel, F., Holvoet, T. and Ferber, J., *Environments for Multiagent Systems State-of-the-Art and Research Challenges*. In: Weyns, D. et al., editors. *Proceedings of the first Workshop on Environments for Multi-Agent Systems (E4MAS 2004)*, Lecture Notes in Artificial Intelligence, volume 3374, Springer Verlag, 2004, pp. 1-47.