Comparison of Reproduction Schemes in an Artificial Society for Cooperative Gathering

A.E. Eiben G.S. Nitschke M.C. Schut

Artificial Intelligence Section
Department of Computer Science
Vrije Universiteit Amsterdam
De Boelelaan 1081a, 1081 HV Amsterdam, NL
gusz@cs.vu.nl nitschke@cs.vu.nl schut@cs.vu.nl

Abstract

This paper compares reproduction schemes for adaptive behavior in an artificial society, where the collective task of the society is the gathering of resources in an artificial environment. The environment is randomly distributed with varying quantities of different resource types, where different resource types yield different fitness rewards for agents that successfully gather them. Gathering of the more valuable resource types (those yielding higher fitness rewards) requires cooperative behavior of varying degrees (a certain number of agents working collectively). We compared reproduction schemes over three dimensions. The first was a comparison of agents that could reproduce only at the end of their lifetimes (*single reproduction at the end of the agent's lifetime*) and agents that could reproduce several times during their lifetime (*multiple reproduction during lifetime*). The second was a comparison of agents that could reproduce only with agents in adjacent positions and agents that could reproduce with agents at any position in the environment. The third compared different methods for deriving the number of offspring produced and the fitness share given to each offspring, as well as stochastic variants of these methods. Results indicate that the single reproduction at the end of the agent's lifetime scheme afforded the artificial society a higher level of performance in its collective task, according to the evaluation criterion, comparative to artificial societies utilizing the multiple reproductions during lifetime reproduction scheme.

1 Introduction

Our research interest can be best described by the term Emergent Collective Intelligence (ECI)[1]. It is rooted in the artificial society simulations field in that it concerns groups of agents, specifically, collectives, which develop certain properties bottom-up. The applications we envision include engineering tasks.

We are interested in the design of cooperative behaviors in groups of agents, where such cooperative behavior could not be developed or specified *a priori*. The key idea is that a desired group behavior emerges from the interaction of the component agents, where no single agent would be able to accomplish the task individually, the task is predefined, and the environment is unknown. The end goal of such an artificial social system would be the transference of a cooperative behavior design methodology to a physical system (for example: multi-robot) that has a specific and well-defined task in an unexplored environment. For example, we envisage the use of such a methodology in swarm-robotics (Nolfi *et al.* 2003) for the gathering of resources in hazardous locations (for example: the surface of Mars or a deep-sea ocean bed). Hence, associating a concrete task with the artificial social system introduces the engineering or design element. If one can measure how well the given task is performed, we have a natural optimization criterion. Consequently, a well-calibrated system will be one where the evolutionary mechanisms (and probably other adaptive features) are able to generate high quality collective behaviors efficiently.

In this paper we consider the task of collective gathering, where a group of agents need to explore their environment in order to find some resources, mine them and collect them at a central location. The formal objective here can be expressed by the total value of resources gathered together in a given amount of time. The system, the environment, and

[1] http://www.cs.vu.nl/ci/eci

the task will be described in *Section 3: Simulator, Environment and Agents*.

As for the agent collective we use an adaptive artificial social system where our technical research goal is to establish what reproduction mechanisms lead to the best results in terms of the total value of resources gathered. In particular, we investigate:

1. Two reproduction schemes, single reproduction at the end of the agent's lifetime (SREL) and multiple reproduction during an agent's lifetime (MRDL)
2. Two mate selection methods locally restricted mating versus panmictic mating.
3. Two methods for determining the initial fitness of new individuals at birth, and for both methods we applied:
   3a. A deterministic variant
   3b. A stochastic variant

These issues will be discussed in *section 4: Experiments* and *section 5: Analysis and Discussion*.

2 Related Literature

This section presents a brief overview of prevalent results pertaining to the study of emergent cooperative behavior, particularly: cooperative gathering and transport, within simulated swarm-based systems. The term swarm-based systems refer to artificial societies containing potentially thousands of agents. Results reviewed maintain particular reference to research that uses biologically inspired design principles and concepts, such as emergence, evolution and self-organization, as a means of deriving cooperative behavior to accomplish tasks that could not otherwise be individually accomplished.

The study of the synthesis of collective behaviour, particularly the emergence of cooperation, is a research field in which there has been little work done in both simulated (Iba, 1996) and real world (Quinn, 2000) problem domains. Traditionally collective behaviour and multi-agent systems have been studied using a top down classical approach. Such approaches have achieved limited success given that it is extremely difficult to specify the mechanisms for cooperation or collective intelligence in all but the simplest problem domains. The investigation of artificial evolution relating to emergent collective behavior, specifically cooperation, remains a relatively unexplored area of research in the cooperative gathering and transport problem domain.

With relatively few exceptions, and then only in multi-robot systems containing relatively few robots (Mataric, 1992), the majority of research in emergent cooperative behavior is restricted to simulated problem domains given the inherent complexity of applying evolutionary design principles to collective behaviors in groups of real robots (Floreano and Nolfi, 2000). This is especially true in swarm-based systems, which by definition contain thousands of individuals.

Within simulated swarm-based systems there has been a significant concentration of research on the study of emergent behavior in artificial ant colonies (Deneubourg *et al.* 1987). Certain artificial life simulators and applications have popularized studies of swarm-based systems. These include *Swarm* (Daniels 1999), *MANTA* (Drogoul *et al.* 1995), *Tierra* (Ray, 2001), and *Avida* (Adami, 1994).

Drogoul *et al.* (1992a; 1992b), (Drogoul and Ferber, 1992) presented a simulation model of social organization in an ant colony termed: *MANTA* (*M*odel of an *ANT*-hill *A*ctivity), which was designed to explore the contribution of emergent functionality such as division of labor on emergent cooperation. Results elucidated that emergent division of labor improved the efficiency of emergent functionality in the population. Such emergent functionality included cooperative foraging and sorting behavior. The authors concluded that the notion of emergent cooperation remains very unclear, difficult to define, and that many of the behaviors viewed as cooperative emerged as a result of the competitive interaction that occurs between individuals in a constrained environment with limited resources.

As part of the swarm-bots initiative, Nolfi *et al.* (2003) conducted several experiments to address the problem of how a group of simulated robots (s-bots) could coordinate their movements and actions so as to cooperatively move objects in the environment as far as possible within a given period of time. Nolfi *et al.* (2003) conducted a set of experiments designed to facilitate emergent cooperative behavior, where a group of eight s-bots were connected to an object, or connected so as to form a closed structure around an object, and were given the task of moving the object as far as possible in the least amount of time. In the first set of experiments the eight s-bots used what the authors termed the *ant formation*, which connected all s-bots to the object, but there were no links between the s-bots themselves. The result was dependent upon the weight of the object, such that the s-bots cooperatively negotiated to either push or pull the object to their destination. In the second set of experiments, s-bots were assembled so as to form a circular structure around the object. The results were similar to those obtained with the ant-formation, with the exception that the s-bot formation deformed its shape so that some s-bots pushed the object, while other s-bots pulled the object. The mechanism deemed to be primarily responsible for these results was the neural controllers of individual s-bots, which evolved the capability to cooperatively coordinate movement when connected to either each other or the object. That is, each s-bot was inclined to follow the direction that the majority of s-bots followed at a given time.

From this overview of these different research efforts, associable by similar tasks and the general research topic of emergent cooperation, it is obvious that some formalization of mechanisms for the design and analysis of emergent cooperation is needed. Specifically, if emergent cooperative behavior in swarm systems was sufficiently understood, purposeful design of cooperative behavior could be applied to benefit a variety of application domains including telecommunications (Di Caro and Dorigo, 1998), space exploration (Brooks and Flynn, 1998) and multi-robot systems (Mitsumoto *et al.* 1995).

3 Environment and Agents

The experiments presented in this paper were performed with our simulation framework: JAWAS[2]. Using this framework we implemented a particular environment and agents populating this environment.

3.1 Swarm-Scape

Swarm-Scape is a specific swarm-based model implemented within the JAWAS simulation framework. Swarm-scape utilizes an initial population of 1000 agents, placed at random positions on a grid-cell environment with a 50 x 50 resolution. A maximum of 4 agents can occupy any given grid-cell within the environment. Also, a home area spanning 4 x 4 grid-cells is randomly placed somewhere within the environment. This home area is where each agent must deliver resources that it is transporting. The process of mining, transporting, and delivering a resource is termed gathering.

Within the environment there exist three types of resources: gold, iron and stone. It is essential in our design that resources also have a *value* that can differ for different types of resources. In particular, in our present system one stone-unit is worth of 1 abstract unit of value, one iron-unit is worth 2, and one gold-unit is worth 4.

Initially, there is some quantity, defined in terms of resource units, of each resource type. For each grid-cell, a maximum quantity (number of resource units) of each resource is specified, and for all grid-cells the re-grow rates (number of resource units that are replenished per simulation iteration) of each resource is specified. Each of these resources has different properties pertaining to its value and cost to transport for each agent.

[2] JAWAS: Java Artificial Worlds and Agent Societies, can be downloaded from http://www.cs.vu.nl/ci/eci/

In order to mine each resource some degree of cooperative behavior is necessitated. Specifically, to mine a unit of gold (the most valuable resource), 4 agents need to be situated on the same grid-cell. To mine a unit of iron (the medium valued resource), at least 3 agents need to be situated on the same grid-cell. To mine a unit of stone (the least valuable resource), only a single agent needs to be situated on the grid-cell. For the purposes of the experiments described within this paper, the term *cooperation* was defined as the instance when at least two agents, situated on the same grid-cell, simultaneously attempted to mine the same resource unit.

3.2 Task Environment

The task of each agent in the environment is the gathering of the highest possible value of resources during the course of its lifetime. This task was interfaced to the agent collective by using the value of the resources gathered by an agent, where gathered value translates into fitness rewards. In our system, fitness was used as a metaphor of energy: performing actions costs fitness units. Furthermore, fitness also played its conventional role in survivor selection: if an agent's fitness reaches zero, it dies.

The particular method we used to reward agents' performance worked as follows. In an instance when a resource unit is delivered to the home area, the agent is given a fitness reward proportional to the total value of the resource units delivered. Specifically, one gold-unit yields a fitness reward of 20 fitness units, 1 iron-unit yields a fitness reward of 10 fitness units, and 1 stone-unit yields a fitness reward of 5 fitness units. The total fitness reward corresponded to the total value of the resources an agent delivered.

The initial amount of gold, iron and stone in the environment was 250, 500, and 1000 respectively, where the number of resource units that could be on any given grid-cell was unlimited. The re-grow rate for each of the three resources was 1 unit per 3 simulation iterations.

3.3 Swarm Agents

Our agents were based on the classical SugarScape design, adopting most of the SugarScape features (Epstein and Axtell, 1996). An agent was able to detect agents and resources for a number of grid-cells determined by a *sight* property. Specifically, an agent was able to detect the number of agents, and the types of resources, in all grid-cells surrounding its current position for a distance (number of cells) given by *sight*.

Each Swarm-Agent used the following set of heuristics in order to determine the action it takes during any given simulation iteration:

```
IF end of life and SREL active THEN reproduce
IF at home THEN unload resources transported
  IF MRDL active THEN reproduce
IF transporting a resource THEN go home
ELSE IF gold detected THEN move to gold
  ELSE IF iron detected THEN move to iron
    ELSE IF stone detected THEN move to stone
      ELSE move to a random cell
```

For any given simulation iteration, each agent was able to move for a number of grid-cells in any position given by the value set for its *move* property. Both the sight and move properties were initially set to one grid-cell. Also, upon initialization each agent was assigned the maximum time for which it would live, assuming that it did not reach zero fitness before this time. This property termed: *death age* was randomly set for each agent to a value between 40 and 80 upon its initialization.

Each agent in the population followed a set of heuristics directing the agent to move, to mine, and then to transport the most valuable resource it could find in the environment. Once an agent had mined as much of a given resource as it could transport (determined by the resource type and the number of units mined), it would immediately begin transporting the resource units back to the home area. Each agent had several properties dictating restrictions on its behavior.

The *maximum gold mining capacity* property specified the maximum number of gold units that each, of 4 cooperating agents, could mine. For these experiments the maximum gold mining capacity property was set to 5. The *maximum iron mining capacity* property specified the maximum number of iron units that each, of at least 3 cooperating agents, could mine. For these experiments the maximum iron mining capacity property was set to 10. The *maximum stone mining capacity* property specified the maximum number of stone units that each agent could mine. For these experiments the maximum stone mining capacity property was set to 20.The *transport-capacity* property determined the maximum number of units of resources a single agent could transport.

An important property for each agent was its *fitness* (that is: the agent's energy rating). At the beginning of each simulation, fitness was randomly initialized for each agent to a value between 90 and 100. Every action taken by the agent cost some portion of its fitness. Mining of any resource type cost one fitness unit. Every grid-cell of distance that an agent moved cost one fitness unit. An agent's fitness could only be replenished when it delivered a resource unit to the home area of the environment.

The initialization settings for each of these parameters is based the most 'appropriate' settings for the given environment, as ascertained in previous experiments (Vink, 2004).

3.4 Reproduction of Swarm Agents

In our system, agents evolved, that is, they underwent variation and selection where the environment performed selection implicitly. Agents with a high fitness (those that performed their tasks most efficiently) were selected for, where as poorly performing agents with not enough fitness died. Variation of agents was accomplished by recombination of agent genotypes.

The core of reproduction was the reproduction cycle where two parent agents created a number of offspring agents via recombining their own genes for *maximum gold mining capacity, maximum iron mining capacity, maximum stone mining capacity* and *transport-capacity* and passing the average of their values onto their offspring.

In this investigation we compared two temporal schemes for reproduction. In the SREL scheme an agent could only perform one *Single Reproduction act at the End of its Lifetime*. That is, when each agent reached the end of its lifetime it selected m mates (partner agents) and then produced a number of offspring according to the particular reproduction method being used. In the MRDL scheme *Multiple Reproduction* acts are executed *During Lifetime*. Using the MRDL scheme, every agent was able to reproduce when a resource quantity was delivered to the home area. Upon delivery of a resource quantity, the agent would receive an immediate fitness reward, and a reproduction cycle would start. During this cycle the agent would select m partner agents from the environment, and then produce a number of offspring according to the reproduction parameters being used.

The second reproduction feature we studied here concerns the spatial distribution of mates for reproduction: *panmictic* versus *locally restricted* mate selection. Using the locally restricted method, an agent could only reproduce with agents in the adjacent grid-cells. In this case, all agents on the same grid-cell or in adjacent grid-cells were taken into account as mates. Using the panmictic method, an agent could reproduce with any other agent anywhere else in the environment. In this case the number of mates m was a random integer between 0 and 10 drawn with a uniform distribution.

Third, we compared two methods for determining the initial fitness given to offspring agents at birth. For both fitness inheritance methods we used a distribution mechanism where 90 percent of a parent agent's fitness was passed onto and divided among its offspring and we divided the total amount of fitness to be inherited (x) over the number of children (n) equally, that is, giving each offspring agent $y = x/n$ fitness units. The parameters to distinguish the investigated methods were n and y.

Using the first method, n, the number of offspring to be produced was predefined and y was derived for each reproduction act by dividing the actual value of x for the two given parent agents by n. In the second method, the fitness share y was predefined and n was determined as x/y (rounded up). The values we used for our experiments are $n = 5$ for the fixed number of offspring method and $y = 10$ for the fixed offspring fitness method.

For both fitness inheritance methods we applied deterministic and stochastic variants. The deterministic variants simply used outcomes of the calculation (rounded up, when needed). The stochastic variants were the same two methods, though random noise was added to the fitness share (in the case of the first method), or random noise to the number of children produced (in the case of the second method). In the case of the first stochastic variant, the random noise was generated within the range between -1 and +1 by a uniform distribution, and in the case of the second variant, random noise was generated within the range of -5 and +5.

4 Experiments and Results

We designed our experiments along three parameter dimensions and two values for each dimension as outlined in the research objectives:

1. Reproduction scheme: SREL versus MRDL.
2. Mate selection method: panmictic versus locally restricted.
3. Fitness inheritance method: fixed n or fixed y.

This led to 8 different experimental setups, although since we also compared a deterministic and a stochastic variant for the inheritance methods, the total number of different experimental setups was 16. For each of them we performed 50 independent runs (using different random initialization parameters), where one run was executed for 2000 iterations.

4.1 Simulation Monitors

Within each simulation, several experimental monitors are set as objective measures for the performance of the society across multiple generations of agents. The first and second are the *number of agents* and the *average value gathered cooperatively* since it is these that determine the value of resources gathered together in a given amount of time, which is our formal objective. The *average fitness* of the population and the *average distance to home*, which describes the population density, are additional measures illuminating details on the overall behavior of the artificial society.

As presented in section 3, cooperative behavior was evaluated according to the total value of each resource: gold, iron, and stone, gathered by the agent population over the course of a given simulation. Specifically, the measure of cooperative behavior is the total value gathered cooperatively, which includes all resource types gathered by the society over the course of the simulation. Submeasures of this are: value of gold gathered cooperatively, value of iron gathered cooperatively, and value of stone gathered cooperatively. These measures can be simply monitored via the GUI and saved for off-line analysis later on, but are not reported in the present paper.

4.2 Results

Figures 1 through to 8 present results attained for the objective measures described above with all 16 different setups. The presentation principle we follow is to use a table style arrangement, with four rows and two columns. Here, each row belongs to one of the measures; the two columns correspond to the two reproduction schemes we investigated. A cell in this table contains a graph divided into a right-hand side and a left-hand side histogram, belonging to the two methods for distributing the parents' fitness over the offspring. Within each histogram deterministic and stochastic variants of these methods are further distinguished by their left/right position. Finally, the two colours are used represent the two mate selection methods.

5 Analysis and Discussion

As mentioned in the introduction, our formal objective is to maximize the total value of resources gathered. To this end, the average value gathered collectively and the average number of agents is essential, as their product indicates how well the population performs.

The reproduction scheme turned out to be one the most influential features in our study, that is, the feature with the highest impact on performance. The impact was most prominent on population sizes. Using the multiple reproductions during lifetime scheme (MRDL) the population sizes varied in a range that was around one tenth of population sizes under the single reproduction at the end of lifetime (SREL) scheme. This is remarkable, in that the number of reproduction cycles was much lower when agents are only allowed to mate once in a lifetime. Apparently, it is worthy to "save" fitness for a longer period and create offspring only in a "rich" state. Perusing the average values gathered one could observe that the impact of the reproduction scheme is much less (as presented in figures 2 and 6). Differences are at most of a factor 2 to 3, sometimes in favour of SREL, sometimes not. Concerning the net effects on total value gathered by the whole population[3] the SREL scheme is the clear winner.

Interestingly, the average fitness values were much less sensitive to these reproduction schemes. In 8 out of the 16 experiments average fitness values did not differ significantly for the SREL and MRDL schemes (as illustrated in figures 3 and 7). In the other 8 cases they did differ in about a factor 3 to 5 in favour of the MRDL scheme. The figures on the average distance to home measure, disclose that the MRDL scheme evolved smaller and denser populations.

The investigated options for the mate selection method, panmictic versus locally restricted reproduction, showed no significant differences in performance for our task environment.

For the inheritance method we could make observations quite similar to those about reproduction schemes. The most affected measure was the population size with differences up to a factor 10. Variations in the average value gathered were much less, up to a maximum of factor 2 to 3. Whether the fixed number of offspring (n) or the fixed offspring fitness (y) method worked better depended on the usage of random noise. For instance, using a fixed y in a deterministic way enabled much larger populations than its stochastic counterpart. However, the fixed n method worked much better in the stochastic variant.

6 Conclusions and Future Work

In this paper we presented an artificial society and a particular task the inhabitants of this society needed to accomplish. This task was the gathering (finding, mining, transporting, and delivering) of certain resources. Resources differed in their difficulty to mine, in that they required a different degree of cooperation to be mined. Resources also differed in their value; that is: the rewards an agent would receive upon delivery were different. Resource mining difficulty and value were related: more difficult resources were worth more.

We investigated reproduction mechanisms within this society and found that two features clearly influenced the performance of the agent population. Firstly, the results of our investigation show that the single reproduction at end of lifetime (SREL) scheme yielded a higher total amount that the population gathered comparative to the multiple reproduction during lifetime (MRDL) scheme. The second feature with a high influence was the fitness inheritance method. The best method depended upon the right combination with either a stochastic or deterministic variant. In particular, we found that the stochastic fixed number of children and deterministic fixed offspring fitness outperformed their counterparts.

The overall best combination of the investigated aspects of the reproduction mechanisms within our world was the SREL reproduction scheme with panmictic mate selection and deterministic fixed offspring fitness. This combination yielded twice the performance (total value gathered cooperatively) of the second best combination.

Three future research objectives have been defined based upon the results presented in this paper. The first is to further investigate the mechanisms that lead to the SREL societies attaining a higher performance (value gathered cooperatively) for the given task, though maintaining a comparable fitness to MRDL societies.

The second is to increase the complexity of the agent controllers and evolutionary process, giving agents the capacity to learn during their lifetimes, as well as evolution the capacity to modify genotypes based upon lifetime behaviors (collective or individual). Modifying the evolutionary process such that a greater part of the agent genotype is subject to evolution would also likely yield greater complexity and diversity in emergent behaviors.

The third is to measure the impact of the number of offspring produced upon the given task. Specifically, to investigate if societies that produce many offspring with small fitness shares have superior performance compared to societies that produce few offspring with relatively large fitness shares.

Forthcoming results will be published on different scientific forums; for locating them conveniently one can visit: http://www.cs.vu.nl/ci/eci.

[3] The total value gathered was the average value gathered (figures 3 and 7) multiplied by the average number of agents (figures 1 and 6).

SREL: Single **R**eproduction at **E**nd of **L**ifetime

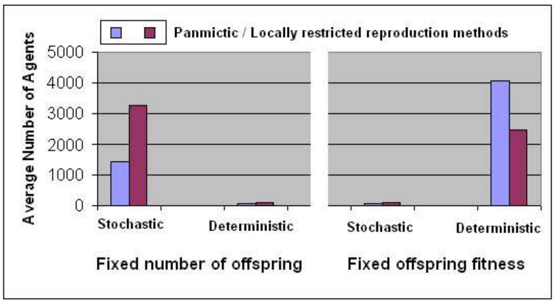

Figure 1: The average number of agents, when using the *SREL* reproduction scheme (Note the scale for the average number of agents in comparisons with figure 5).

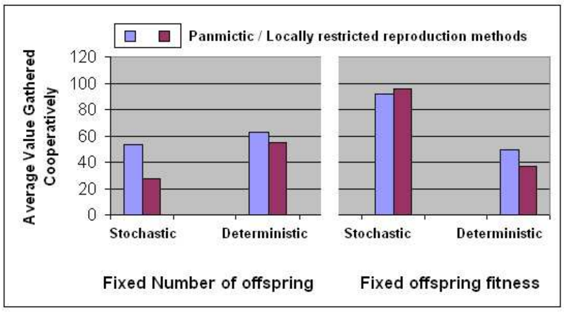

Figure 2: The average resource value gathered cooperatively by the agent population, when using the *SREL* reproduction scheme.

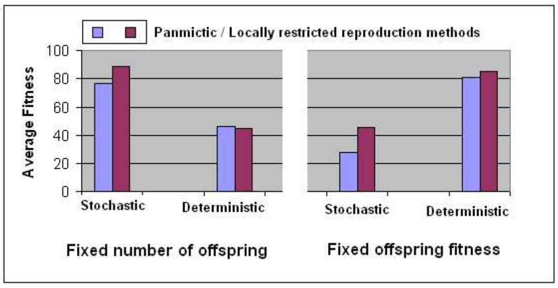

Figure 3: The average fitness of the agent population attained under the *SREL* reproduction scheme.

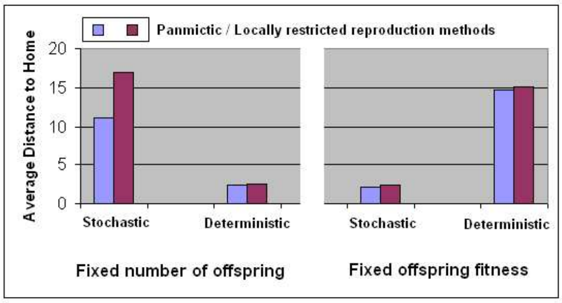

Figure 4: The average distance to home for the agent population, when using the *SREL* reproduction scheme (Note the scale for the average distance to home in comparisons with figure 8).

MRDL: Multiple **R**eproductions **D**uring **L**ifetime

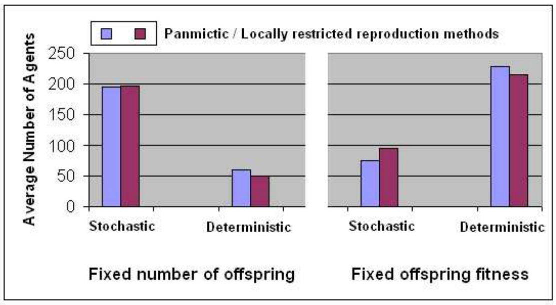

Figure 5: The average number of agents, when using the *MRDL* reproduction scheme (Note the scale for the average number of agents in comparisons with figure 1).

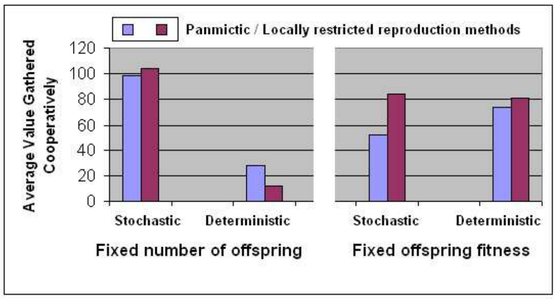

Figure 6: The average resource value gathered cooperatively by the agent population, when using the *MRDL* reproduction scheme.

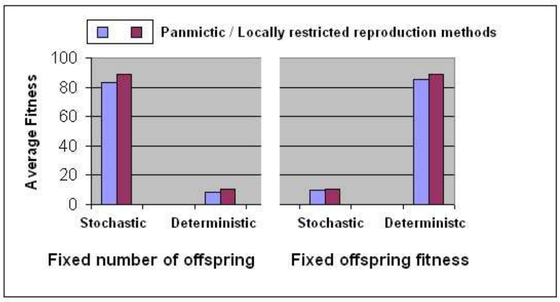

Figure 7: The average fitness of the agent population attained under the *MRDL* reproduction scheme.

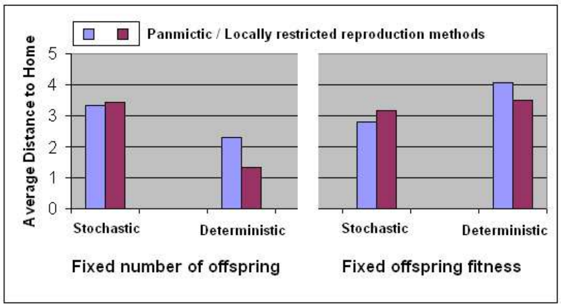

Figure 8: The average distance to home for the agent population, when using the *MRDL* reproduction scheme (Note the scale for the average distance to home in comparisons with figure 4).

Acknowledgements

We would like to thank Nico Vink, Volker Nannen, and Robert Griffioen for their contributions to the JAWAS platform and the inspiring discussions.

References

Adami Chris, Evolutionary Learning in the 2D Artificial Life System Avida, *Technical report*, Kellogg Radiation Lab, Caltech, Pasadena, 1994.

Brooks Rodney and Flynn Anita. Fast, Cheap and Out of Control: A Robot Invasion of the Solar System. *British Interplanetary Society Journal*, 1(1): 478-485, 1989.

Deneubourg Jean-Louis, Goss Simon, Pasteels Jacques, Fresneau Dominique, and Lachaud Jean-Paul. Self-Organization Mechanisms in Ant Societies (II): Learning in Foraging and Division of Labor. *Behavior in Social Insects*, 54(1): 177-196, 1987.

Di Caro Gianni and Dorigo Marco. AntNet: Distributed Stigmergic Control for Communications Networks. *Journal of Artificial Intelligence Research*, 9(1) 317-365, 1998.

Drogoul Alexis, Ferber Jacques, Corbara Bruno, Fresneau Dominique, A behavioral simulation model for the study of emergent social structures, In *Toward a practice of autonomous systems*, Varela Francisco and Bourgine Paul, Eds, Kluwer Academic Publishers, Brussels, pp. 161-170. 1992.

Drogoul Alexis, Corbara Bruno, Fresneau Dominique. Applying Etho-modeling to social organization in ants, In *Biology and Evolution of Social Insects*, Billen Johan, Ed, Leuven University Press, Leuven, pp. 375-383, 1992.

Drogoul Alexis and Ferber Jacques. Using reactive multi-agent systems in simulation and problem solving, In *Distributed Artificial Intelligence: theory and praxis*, Avouris Nicholas and Gasser Les, Eds, Kluwer Academic Publishers, Brussels, pp. 53-80. 1992.

Drogoul Alexis, Corbara Bruno, Lalande Steffen, MANTA: New Experimental Results on the Emergence of (Artificial) Ant Societies, In *Artificial Societies: the computer simulation social life*, Gilbert Nigel and Conte Rosaria, Eds, University College of London Press, London, pp. 190-211, 1995.

Epstein Joshua and Axtell Robert. *Growing Artificial Societies: Social Science From The Bottom Up*, Brookings Institute Press, 1996.

Eiben A.E and Smith Jim. *Introduction to Evolutionary Computing*, Springer, 2003.

Floreano Dario and Nolfi Stefano. *Evolutionary Robotics*. MIT Press, Cambridge, 2000.

Iba Hitoshi. Emergent cooperation for multiple agents using genetic programming. In *Parallel Problem Solving from Nature, PPSN IV.* LNCS 1141, Springer-Verlag, Berlin, pp. 32-41, 1996.

Mataric Maja. Designing emergent behaviors: From local interactions to collective intelligence, In *Proceedings of the Second International Conference on Simulation of Adaptive Behavior*, MIT Press, Cambridge, pp. 432-441, 1992.

Mitsumoto Naoki, Fukuda Toshio, Shimojima Koji, Ogawa Akio. Micro Autonomous Robotic System and Biologically Inspired Swarm Strategy as a Multi Agent Robotic System, In *Proceedings of the IEEE International Conference on Robotics and Automation*, IEEE Press, Tokyo, pp. 2187-2192, 1995.

Nolfi Stefano, Baldassarre Gianluca, and Parisi Domenico. Evolution of collective behavior in a team of physically linked robots, In *Applications of Evolutionary Computing*, Gunther Raidl, Guillot Agnès, Meyer Jean-Arcady, Eds, Springer Verlag, Heidelberg, pp. 581-592, 2003.

Quinn Mathew. Evolving cooperative homogeneous multi-robot teams. In *Proceedings of the International Conference on Intelligent Robots and Systems*, IEEE Press, Takamatsu, pp. 1798-1803, 2000.

Ray Tom. Overview of Tierra, *Technical report*, Advanced Telecommunications Research Institute: Human Information Processing Research Laboratories, Kyoto, 2001.

Vink Nico. Exploring Communication Dynamics in VU-Scape, Technical report, Department of Computer Science, Vrije Universiteit, Amsterdam, The Netherlands, 2004.