ELSEVIER

BioSystems 66 (2002) 21–30

Bio Systems

www.elsevier.com/locate/biosystems

Two-level evolution of foraging agent communities

Manuel Alfonseca *, Juan de Lara

Department Ingeniería Informática, Universidad Autónoma de Madrid, Ctra. De Colmenar, km. 15, 28049 Madrid, Spain

Received 19 February 2002; received in revised form 29 March 2002; accepted 19 April 2002

Abstract

This paper presents simulation results of artificial foraging agent communities. The goal of each agent in the community is to find food. Once a food source is found, agents eat portions of it and carry some other portions to the nest (in a manner similar to ants) until the food is depleted. Agents may also communicate food positions when they are near each other. They are given a set of genes that control several characteristics, such as their activity, memory, scepticism, lying, etc. These genes are recombined and propagated by means of sexual reproduction. When one nest is superpopulated with agents, it can break in two nests. Agents can communicate only with those belonging to the same nest, which gives rise to emergent situations of competition and cooperation between the agents in the same nest, as well as competition between different nests. Other emergent phenomena such as the propagation of rumours are also studied. © 2002 Elsevier Science Ireland Ltd. All rights reserved.

Keywords: Agent-based simulation; Evolution; Competition; Evolutionary stable strategies

1. Introduction

Computer modelling of large individual communities (Resnick, 1999) is an active area of research, due to the increasing capabilities of today's computers. Modelling is done mostly using cellular automata (Hegselmann and Flache, 1998) or agent-based techniques. The latter formalism is a useful and natural way to conduct complex simulation experiments, in which many autonomous and interacting entities participate. The key abstraction in this methodology is the autonomous agent. According to (Jennings et al., 1998), an agent is 'a computer system, situated in some environment, that is capable of flexible autonomous action in order to meet its design objectives'. Agents interact via discrete events. These simulations are difficult to express in other formalisms.

The concept of evolutionary stable strategies (Maynard Smith, 1982) is based on results from game theory (Von Neumann and Morgenstern, 1953). Within the framework of natural selection, individuals adopt different strategies with frequencies that reflect their success for applying these strategies. For example, some species or societies

* Corresponding author

E-mail addresses: manuel.alfonseca@ii.uam.es (M. Alfonseca), juan.lara@ii.uam.es (J. de Lara).

0303-2647/02/$ - see front matter © 2002 Elsevier Science Ireland Ltd. All rights reserved.
PII: S0303-2647(02)00018-7

may exhibit a truthful behaviour in situations where lying would give immediate benefit. The truthful behaviour occurs because if the lying strategy were adopted by a large number of individuals, the whole group would be affected detrimentally. Thus, the genetic basis for lying has been 'deselected' by evolution.

There are many systems that simulate agent communities which are similar to ours. Some of them have a biological inspiration (Drogoul, 1995; Guérin et al., 1998) and carry out realistic simulations of animal behaviour, such as ants, which communicate by dropping pheromones on the ground. In others, such as (Anderson et al., 1997), the focus is on foraging behaviour and the number of agents is kept low (100 ants). This model uses a modification of the Ollason model (Ollason, 1980, 1987) of hunting by expectation.

Our approach and objectives are different from existing agent simulation systems in the sense that we do not try to simulate the real behaviour of any species or society, but rather experiment with emergent behaviour in a virtual agent colony. Another difference is that our agents communicate directly, rather than by intermediate means (such as dropping pheromones). In this way, we can study the propagation of knowledge in the population. In previous work (Alfonseca and de Lara, 2002), we experimented with similar agents, but with only one nest (so competition between nests was not present). The model we present in this paper uses a genetic algorithm (Holland, 1975), which acts on a genome that controls some aspects of the agent behaviour that is interesting for our purposes, such as lying, scepticism, talkativeness, etc. Emergent behaviour has been observed in these characteristics; the results of the simulation suggest that some combinations of genes lead to non-evolutionary stable strategies.

The paper is organised as follows: Section 2 explains the entities that we have implemented in the simulation, Section 3 shows the properties of our agents, Section 4 explains the behaviour of agents, Section 5 details nest properties and behaviour, Section 6 describes the experiments performed and their results, and Section 7 presents the conclusions and the future work.

2. Simulation entities

Our simulation takes place in a territory, which is a rectangular space where the agents move and the food sources and agent nests are located. A territory is divided into a number of integer space steps and has an associated co-ordinate system. The territory is, thus, quantified and the co-ordinates of a location are always integer.

There are a number of food sources in the territory. They are immobile objects with a certain capacity. Agents will break unit pieces of the food to bring to their nest. When the food capacity of the source is exhausted, the food source disappears and another food source appears spontaneously in a different location, with a random initial capacity chosen from a uniform distribution between 1 and 200. The number of simultaneous food sources is a parameter of the program.

Each agent belongs to a nest. Nests are immobile objects associated to a certain number of agents and contain a certain amount of food, which is used for reproduction purposes.

Agents move around in search of food. When they find it, they break off a unit piece and bring it to their nest. Agents may remember the position of the last place where they found food, communicate that position to other agents belonging to the same nest, rob agents from different nests that possess food, or kill them and use them as food. They can also reproduce while they stay in the nest, assuming there is sufficient food in the nest.

3. Agent properties

One agent is an object that belongs to a particular nest, and is situated in a certain position of the territory. Each agent knows the position of its nest and a maximum of one food location.

Agents have a life expectation, which is given an initial value when the agent is born, and is obtained randomly from a Gaussian distribution with an average of 250 time steps and a S.D. of 100. An agent's life expectation is not a constant, it may increase or decrease depending on the agent's history. When this expectation reaches zero, the agent dies.

If an agent does not know the location of a food position, it explores in a random way. We assign each agent a constant that regulates the number of time steps after which the agent tires of seeking in vain. This is a positive constant for each agent, obtained randomly at birth from a Gaussian distribution with an average of 100 time steps and a S.D. of 25. When agents come back to the nest, they rest there for a while. The number of time steps the agent will rest is computed each time the agent arrives to its nest as a uniform random integer between 1 and 8.

Each agent in the simulation contains a virtual genome with six genes that define several phenotypic characteristics. The first is a lying gene (two bits, or a 0–3 integer), which regulates the agent behaviour when communicating food locations to another agent in the same nest. If the value of this gene is 0, the agent conveys truthful information about the location of the food source. If it is 1, however, the agent conveys that information in an approximate manner (possibly exact) based on location. If the value of the gene is 2, the agent communicates the food source location in a random manner; if it is 3, the agent sends other agents from the same nest to a location symmetrical with respect to the centre of the territory.

The second gene (a scepticism gene defined by four bits, or a 0–15 integer), regulates the extent to which the agent believes the information transmitted by another agent in the same nest, specifically regarding food source location. A value of 0 in this gene means ‘total belief’ in the information, whereas a value of 15 represents ‘total scepticism.’

A third gene (talkativity gene with four bits, or a 0–15 integer), regulates to which extent an agent is ready to communicate the location of a food source to another agent in the same nest. A value of 0 in this gene means that the agent never passes the information, a value of 15 means that the agent speaks always. This gene is independent on the actual information passed, which depends on the lying gene.

The fourth gene (activity gene, two bits or a 0–3 integer), regulates the speed of movement of the agent. A value of 3 in this gene means that the agent moves always; if it is 2, it moves 75% of the time steps; if it is 1, it moves 50% of the time; if it is 0, it moves 25% of the time.

The forgetfulness gene (two bits, or a 0–3 integer), regulates the probability that an agent forgets the location of the last food source it found or was told about. A value of 0 in this gene means that the agent never forgets. A value of 3 corresponds to a 0.25 probability of forgetting.

Finally, the aggressivity gene (four bits, or a 0–15 integer), regulates the aggressivity of an agent against agents in different nests. If an agent with no food meets another agent bringing food to a different nest, it may rob the other agent if its aggressivity is at least two units larger than the other’s. If an agent with no food meets another agent with no food from a different nest, the first may kill the second and use it as food if the former’s aggressivity is at least eight units larger than the other’s.

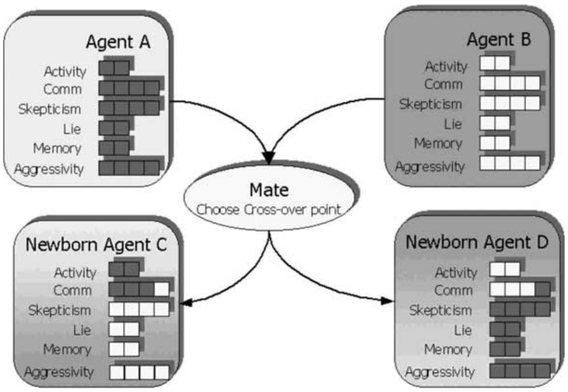

Fig. 1. Agents’ reproduction.

4. Agent behaviour

If the agent is outside the nest, before it moves to a different location, the activity gene is checked. If the agent moves and is bringing food to its nest, it takes the shortest path. If it reaches the nest, it leaves the food there, and as a prize receives an increase of 50 time steps in its life expectation and prepares to rest in the nest for a random number of time steps. When the rest time has elapsed, the agent leaves the nest. The current time step is remembered and will be used to compute the time

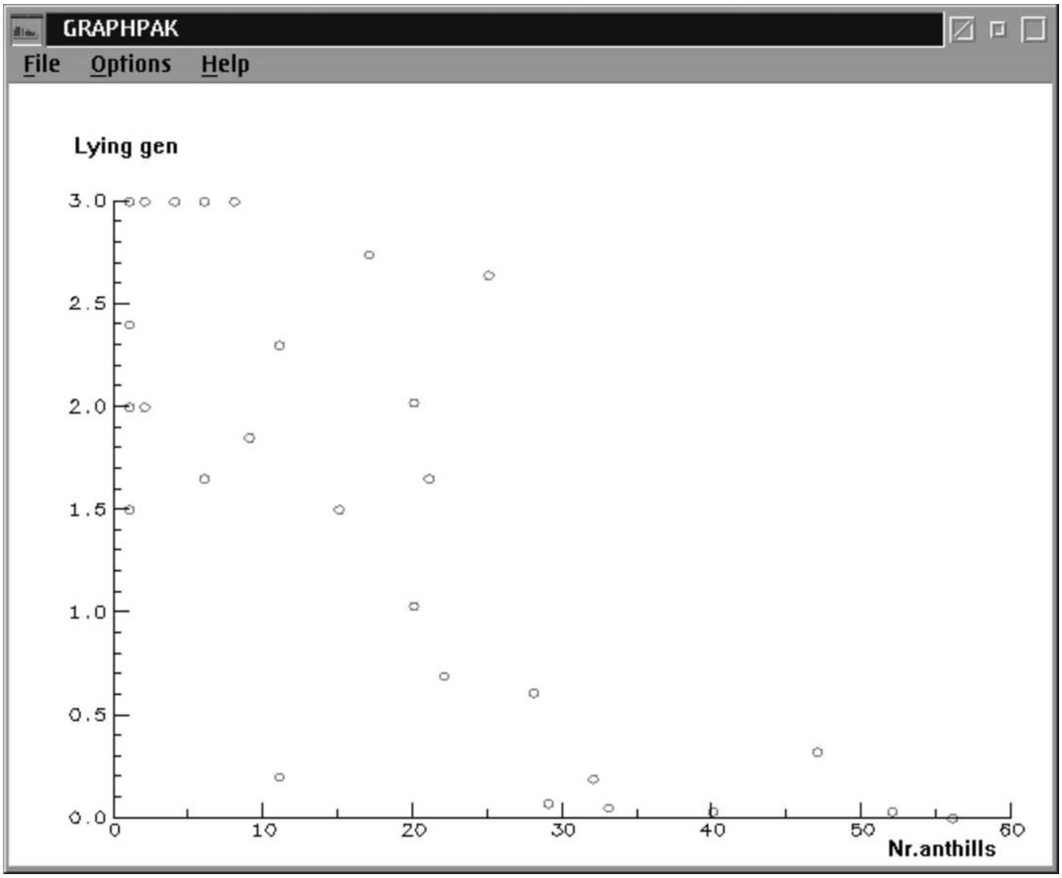

Fig. 2. Lying gene evolution as a function of the number of nests generated.

step when the agent will get tired of searching food in vain and will return to the nest by the shortest path.

If the agent has no food and is out of its nest, but knows where a food source may be, it moves toward that location by the shortest path. If it reaches that location and there is no food there, it resets the switch indicating that it knows the location of food and starts looking around. Forgetful agents may forget (randomly, according to their forgetfulness gene) the position where they think there is food. The agent may also be outside the nest without knowing any food position. In this case, it moves in a random direction, except towards the nest.

If the agent finds food in its new position (whatever it is) it picks a piece of food to bring to its nest, depleting the food source, remembers where it is, sets the switches indicating that it brings food and knows where it is, and its current life expectation is increased by 50.

Agents may also speak to other agents from the same nest. Both agents must be alive and in the same location, outside the nest. The first agent must know (or believe it knows) where a food source is, while the other agent must not know it. Depending on its talkativity gene, the first agent may refuse to speak. Depending on its scepticism gene, the other agent may refuse to listen. The actual location communicated depends on the lying gene of the first agent. If the agents communicate, the agent that receives information pays the speaking agent five time steps from its own life expectation, regardless of whether it is being told the truth. Agents with false information can then spread it. Thus, lies can cause the appearance of rumours in the population, that is, a false information about a food location that is propagated to a large number of members. Following a rumour can be harmful to the agents, because it is a waste of energy. Scepticism, and to some extent a low talkativity, are means to protect agents against lies; one would expect that, in environments with a high number of lying agents, scepticism would rise, while it would decrease in truth-dominated environments.

An agent can reproduce with another agent from the same nest. Both agents must be alive and resting in the nest. The number of agents associated to this nest must also be smaller than the maximum allowed, and the nest must contain at least 100 food units. If all those conditions hold, two new agents are generated, with the same genomes as their parents, although two genetic operations are applied: crossing over (both genomes are split at the same random position and rebuilt by crossing the parents' information) and mutation (one bit is switched in the genome of one of the offspring with a 0.005 probability).

Fig. 1 shows a scheme of the agents' reproduction.

5. Nest properties and behaviour

A nest is an object situated in a territory. For each nest, it is necessary to record the amount of food currently in the nest, because mating in the nest will depend on this quantity.

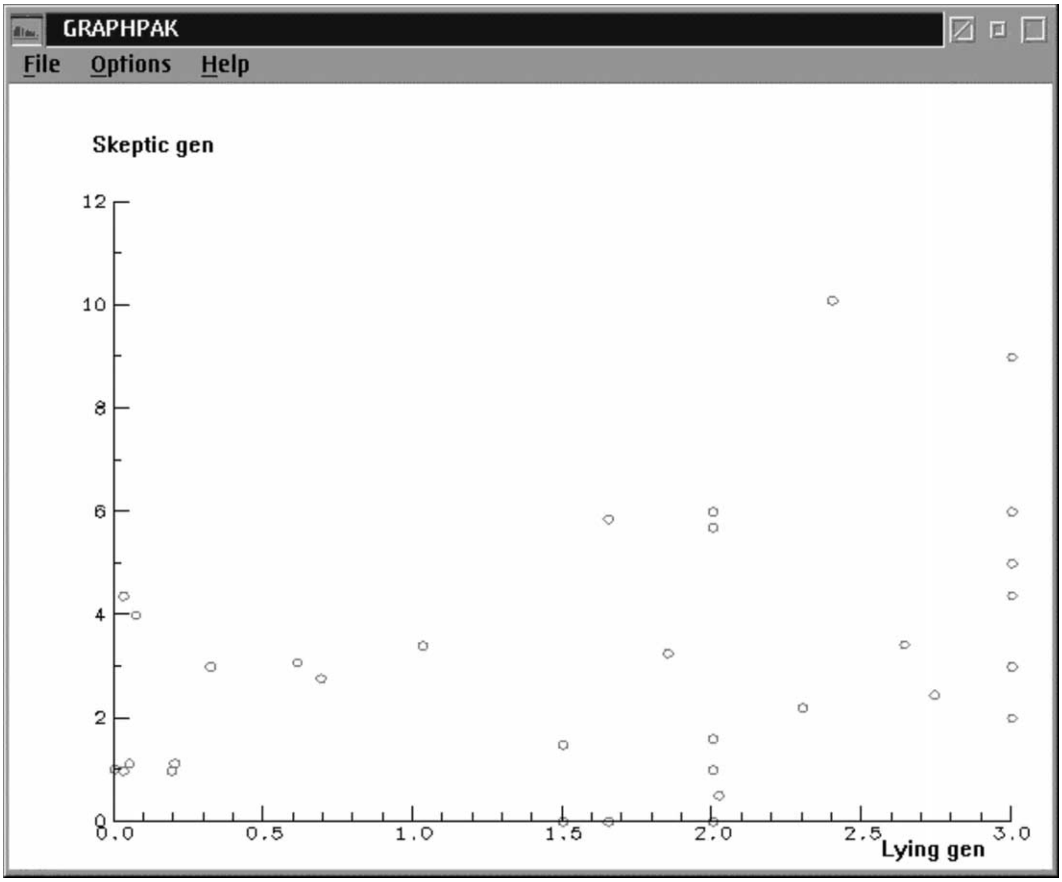

Fig. 3. Sceptic gene evolution as a function of the lying gene.

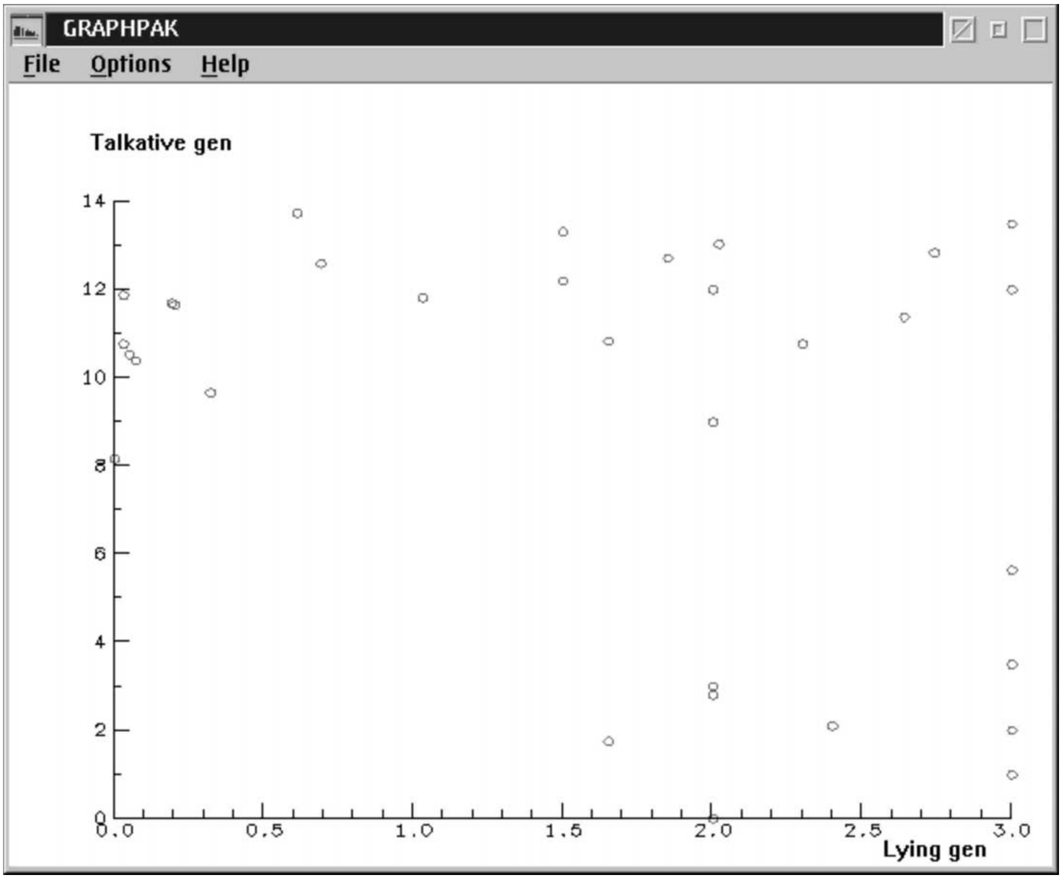

Fig. 4. Talkative gene evolution as a function of the lying gene.

Each nest has a list of its agents and is responsible of moving them, making them speak to one another, and making them reproduce if the appropriate conditions hold. Nests also hold the number of living agents and the maximum number of associated agents. They split into two different nests if the number of living agents is equal to the maximum and the amount of food in the nest is larger than a given value. A new nest is created in a random position in the territory, and half the agents in the old nest are transferred to the new one, together with half the food. The two nests become independent and their agents become hostile to one another. The maximum number of simultaneous nests in the territory is a parameter of the program. If that maximum is reached, nest splitting is inhibited, even though the proper conditions hold, until one of the current nests disappears. A nest dies when the number of its living agents decreases below five.

6. Experiments performed

Thirty-two experiments were performed with the following conditions: the territory is a 50 by 50 rectangle. A single nest is created at the beginning of the simulation, containing 100 living agents, with a maximum of 200 agents per nest and 500 initial pieces of food. A certain number of food sources, chosen for each simulation run, are distributed randomly throughout the territory. Each experiment runs for 50 000 time steps or until all nests disappear. All the experiments aborting before 10 000 time steps were discarded and are not included in the 32. Four experiments that aborted before 20 000 time steps, have been considered in the results analysis. Twenty-eight experiments reached a stable situation and were interrupted after 50 000 time steps, except two, which were allowed to run to 150 000 time steps, to study further developments.

We analysed the 32 experiments and found that the activity gene was selected quickly, so that only the most active agents (agents with the maximum value of this gene) survived after a few thousand simulation steps. A similar effect applied to the forgetfulness gene, which was selected so that only the agents that never forgot (those with the minimum value of this gene) survived after a few thousand steps. These two effects occurred in every experiment.

The effects of evolution on the lying gene were more complicated: two different attractor situations developed. In one, usually associated with food scarcity, the initial nest never split, or very few nests appeared, and the lying agents (with this gene valued at 2 or 3) were positively selected. This happened in half the experiments (16), including the four prematurely aborted cases. The average number of nests generated per experiment was six. The average final value of the lying gene after 50 000 steps was a number between 2 and 3.

In the second situation, usually associated with food abundance, the initial nest splits many times and truthful agents (with the gene valued at 0) were positively selected. This happened in ten experiments. The average number of nests generated per experiment in this case was 35. The average final value of the lying gene after 50 000 steps was very near 0.

There was an intermediate situation, where the lying gene remained in an intermediate value (between 1 and 2) which meant that truthful and lying agents coexisted in the same nest in a possibly unstable equilibrium. This happened in six experiments (see Fig. 2).

This effect may be explained by a two-level evolution effect. When food is scarce, there is not enough food for the nest to split frequently, only a few nests share the territory, and competition inside a single nest dominates, that is, agents compete mostly with other agents in their own nest. In this situation, lying gives agents a competitive advantage: when they find a food source, they may come back again and increase their own life expectation against that of their nest mates by lying about the location of the source.

When food is abundant, however, nests split frequently and agents compete mainly with agents from different nests. In this case, telling the truth to their mates gives the nest as a whole a competitive advantage against other nests, so that nests dominated by lying agents are probabilistically eliminated.

The evolution of the gene for scepticism is less marked, although there is a difference between the two attractor situations mentioned above. When truthful agents dominate, they also tend to be less sceptical (with values of this gene in the 0–4 interval). When lying agents dominate, however, more sceptical agents may develop, but this may not be the case (average values of the gene are spread throughout the 0–10 interval). See Fig. 3, where it is clear that a combination of truthful and highly sceptic agents is not an evolutionary stable strategy.

A similar situation affects the gene for talkativity, which reaches high values (in the interval 10–14) when truthful agents dominate, while it is spread throughout the value interval (0–14) when lying agents dominate. See Fig. 4.

Finally, the gene for aggressivity is also related to the number of nests generated by the experiment. If the number is 8 or greater, aggressivity against agents in a different nest is strongly selected, and the average value of this gene tends to a high value (around 14). This means that nests whose agents are not aggressive die, because they are at a disadvantage with respect to nests with aggressive agents.

If there are only a few nests, however, the value of this gene may end wherever in the 6–14 interval. See Fig. 5.

7. Implementation

The model was implemented progressively, by first using a prototype written in APL2, where not all the functionality described in the paper was implemented.

Once we were familiar with the model and made some refinements, we implemented a full version in our object-oriented declarative simulation language OOCSMP (Alfonseca and de Lara, 2000) and translated it into C++ and JAVA by means of our C-OOL compiler. Fig. 6 shows a moment in the execution of a JAVA simulation program generated by our compiler.

The front window shows a map with the agents and food locations. The agents knowing a food location are shown in lighter colour. The nest is shown in the middle of this panel. On the right panel in the same window we can observe a graphic of the number of agents (upper curve, in green) and the number of agents knowing a food location (lower curve, in black). Note that some of these agents may be mistaken due to rumours. The background window shows a listing of some simulation data: the simulated time and the average of the activity gene values are visible.

Finally, and for the purpose of maximum efficiency, another full version was written directly in C++ using the object-oriented classes described in the previous sections. This version was more efficient than the OOCSMP version, although it was less compact and less readable.

8. Conclusions and future work

Two-level evolution emerges as a spontaneous result of the aggregation of agents in groups

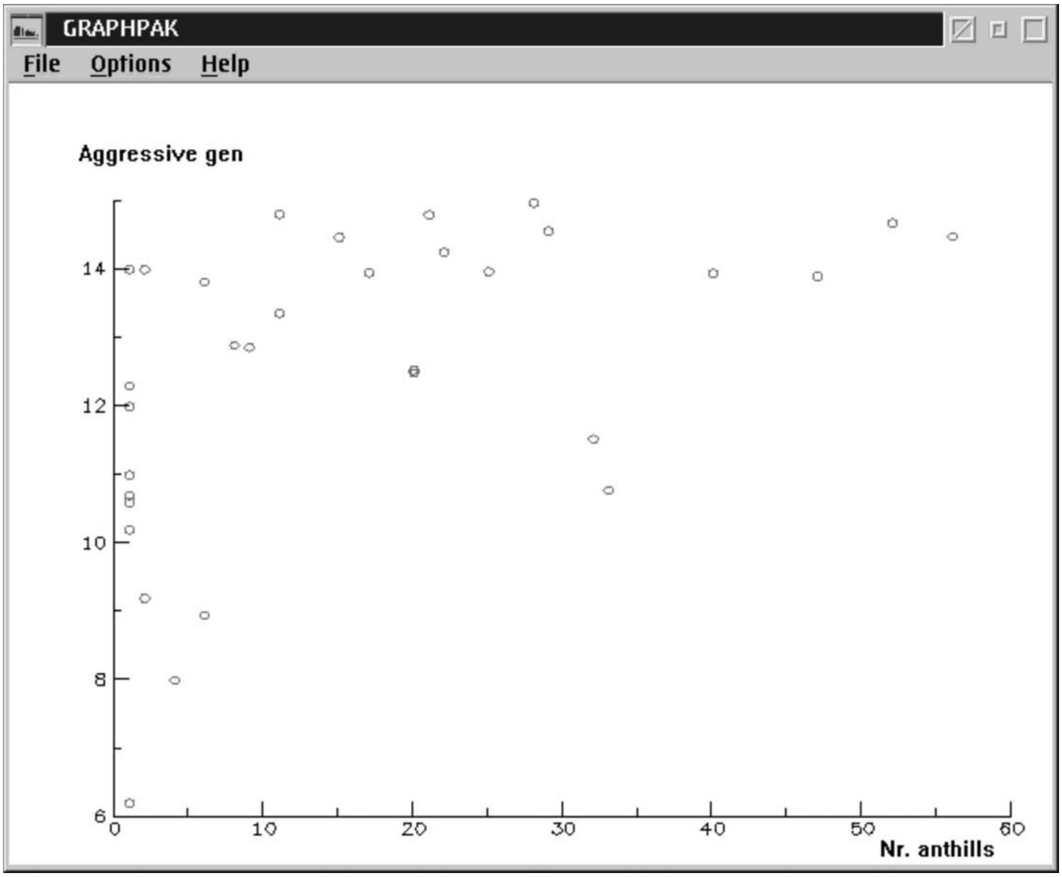

Fig. 5. Aggressivity evolution as a function of the number of nests generated.

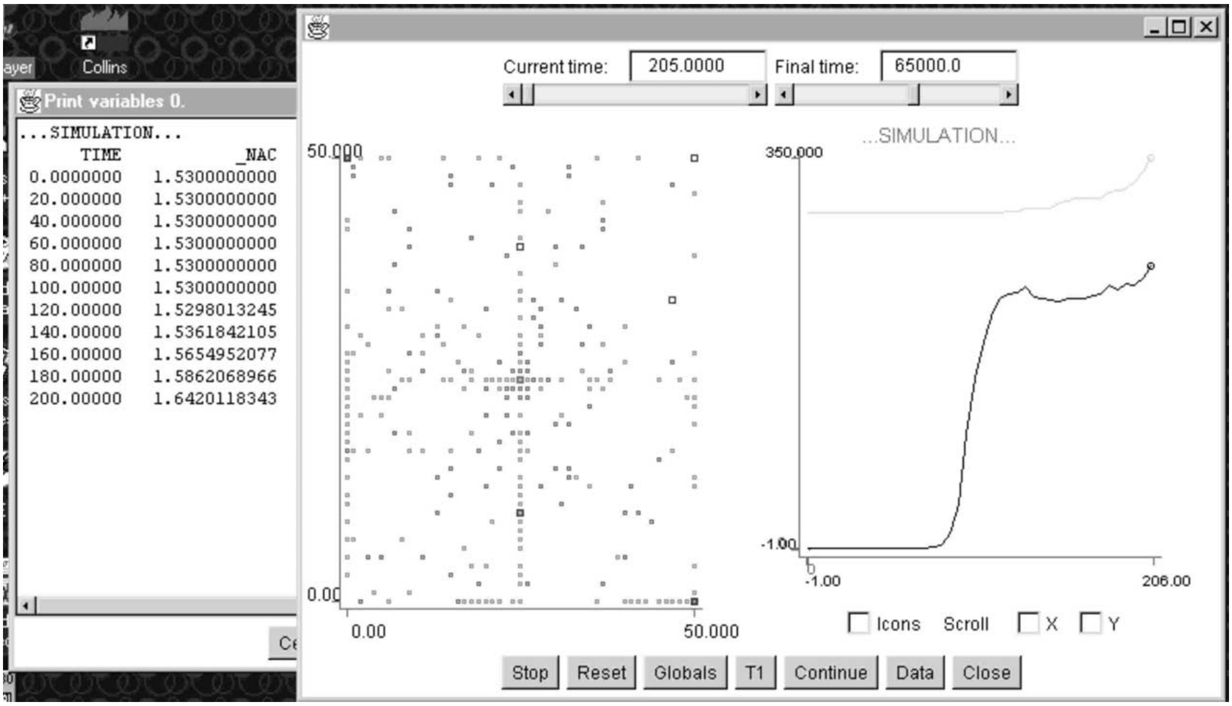

Fig. 6. A JAVA simulation program generated by C-OOL.

(nests). Different results are obtained when the number of aggregations generated is small or large. This depends on whether the food is scarce or plentiful. In the first case, agents compete mainly with other agents in the same aggregation; in the second, they compete mainly against agents in different aggregations, and cooperate with agents in the same aggregation. Great variations in the lying gene are observed depending on these situations.

Another emergent phenomenon is the appearance of rumours. This can happen due to two factors, one natural and another caused by the behaviour of agents.

- Depletion of a food source. In this situation some agents may believe that the food is still there and propagate false information. This may cause a slight raise of scepticism in the population, even when the lying gene has low values.
- Lies. Lying agents have some advantage in environments where truthful and credulous agents dominate.

The models described in this paper do not try to represent the behaviour of any real biological species or society.

In the future, we intend to test more complex situations, such as a territory with obstacles that divide it in several almost isolated sections, to find whether divergent evolution and speciation appears. We shall also test different strategies, such as periodically varying food sources.

It would also be interesting to integrate these experiments with our research in communication between agents (de Lara and Alfonseca, 2000). Agents in each nest would be given their own vocabulary for their spatial movements. Agents belonging to different nests would not be able to reach an understanding. When one nest breaks in two, both nests would initially share the same vocabulary, but the vocabularies would change with time due to interaction.

Acknowledgements

This paper has been sponsored by the Spanish Interdepartmental Commission of Science and Technology (CICYT), project numbers TEL1999-0181 and TIC2001-0685-C02-01.

References

Alfonseca, M., de Lara, J., 2000. Automatic generation of Simulation-based Web courses and model documentation. Transactions of the Society for Computer Simulation 17 (3), 106–119.

Alfonseca, M., de Lara, J., 2002. 'Simulating evolutionary agent communities with OOCSMP'. Proceedings of the ACM Symposium on Applied Computing SAC'02, AI Track, March 2002, Madrid, Spain, pp. 11–15.

Anderson, C., Blackwell, P.G., Cannings, C., 1997. Stochastic simulation of ants that forage by expectation. In: Husbands, P., Harvey, I. (Eds.), Fourth European Conference on Artificial Life. MIT Press, London.

de Lara, J., Alfonseca, M., 2000. Some strategies for the simulation of vocabulary agreement in multi-agent communities. Journal of Artificial Societies and Social Simulation 3, 4 < http://www.soc.surrey.ac.uk/JASSI/3/4/2.html > .

Drogoul, A., 1995. When ants play chess (or can strategies emerge from tactical behaviors). In: Castelfranchi, C., Müller, J.P. (Eds.), From Reaction to Cognition, Lecture Notes in AI, No 957. Springer, Berlin, Heidelberg, pp. 13–27.

Guérin, S., Snyers, D., Fourcassier, V., Théraulaz, G., 1998. Modeling Ant Foraging Strategies by Computer Simulation. Ants, Brussels, Belgium, p. 98.

Hegselmann, R., Flache, A., 1998. Understanding complex social dynamics: a plea for cellular automata based modelling. Journal of Artificial Societies and Social Simulation 1, 3Internet at: < http://www.soc.surrey.ac.uk/JASSI/1/3/1.html > .

Holland, J.H., 1975. Adaptation in Natural and Artificial Systems. University of Michigan Press.

Jennings, N.R., Sycara, K., Wooldridge, M., 1998. A roadmap of agent research and development. In: Autonomous Agents and Multi-Agent Systems. Kluwer Academic Publishers.

Maynard Smith, J., 1982. Evolution and the Theory of Games. Cambridge University Press.

Ollason, J.G., 1980. Learning to forage—optimally. Theoretical Population Biology 18, 44–56.

Ollason, J.G., 1987. Learning to forage in a regenerating patchy environment: can it fail to be optimal. Theoretical Population Biology 31, 13–32.

Resnick, M., 1999. Turtles, Termites, and Traffic Jams, Explorations In Massively Parallel Microworlds, fifth ed. MIT Press.

Von Neumann, J., Morgenstern, O., 1953. Theory of Games and Economic Behaviour. Princeton University Press.