1

Aggressive Signaling Meets Adaptive Receiving: further experiments in Synthetic Behavioural Ecology

Peter de Bourcier and Michael Wheeler
School of Cognitive and Computing Sciences,
University of Sussex, Brighton BN1 9QH, U.K.
`Telephone:+44 1273 678524`
`Fax:+44 1273 671320`
`E-Mail:peterdb@cogs.susx.ac.uk, michaelw@cogs.susx.ac.uk`

Abstract

This paper presents the results of our most recent investigations into the relationship between (i) aggressive communication systems, and (ii) the ecological contexts in which those systems evolve. We perform experiments in a simple synthetic ecology, in which simulated animals (animats) are in competition over food. In the first experiment, each individual has a signaling strategy, which is determined by a form of artificial evolution in which there is no explicit fitness function. We vary the cost of producing aggressive signals, and analyse the resulting population dynamics. The results indicate that the general logic of the handicap principle (according to which higher costs enforce reliability) can apply in the sort of ecological context not easily studied by formal models in theoretical biology. We then extend the model to take account of the fact that the behaviour of an animal in reaction to an incoming signal will be the result not of the signal alone, but also of the degree of importance that that receiver gives to that signal. So, to investigate the effect of receiver tactics on aggressive communication systems, we introduce the concurrent evolution of individual signaling and receiving strategies. Variations in the cost of signaling result in notably different population dynamics. We discuss certain key features of these dynamics, with particular reference to their effect on the reliability of the signaling system.

Keywords: aggression, animal signaling, behavioural ecology, communication.

---

[*] Submitted to the Third European Conference on Artificial Life, Granada, Spain, 4-6 June, 1995.

1 Introduction

When animals of the same species come into conflict, the incidence of unrestrained battles is relatively low. Confrontations tend to revolve around ritualized signaling displays which, more often than not, allow those concerned to conclude matters without the need for actual physical combat. In this paper, we describe the latest extensions to our ongoing work on aggressive signaling (de Bourcier & Wheeler, 1994; Wheeler & de Bourcier, 1994).

Our study takes place within a theoretical framework that we call *Synthetic Behavioural Ecology* ('SBE'). In the biological sciences, behavioural ecology is the discipline that aims to explain *why,* rather than *how,* animals behave as they do (see Krebs & Davies, 1987). That is, behavioural ecologists try to identify the functional role that particular behaviours play in contributing to the survival and reproductive prospects (the Darwinian fitness) of an animal. The SBE-methodology (in common with much of A-Life) is to construct simple synthetic ecologies. By using information from the biological sciences to guide this process, the synthetic environments employed should be simplified and idealized, but not trivial. We then carry out experiments in these synthetic ecologies, with the specific goal of making a contribution to the scientific understanding of how ecological context influences the adaptive consequences of behaviour.

The primary aim of SBE is to contribute to ongoing work in the biological sciences, by providing a distinctive theoretical platform for testing hypotheses about the functional aspects of animal behaviour. SBE may be able to play such a role, because it is pitched at a level in between abstract mathematical models and naturally occurring ecological contexts. However, we must stress that there is no suggestion that SBE provides any easy answers to the difficult problems faced by biologists in this area.[1]

Two SBE-experiments are described. The first is a consolidation of our earlier work on Zahavi's handicap principle (see de Bourcier & Wheeler, 1994). The handicap principle (Zahavi, 1975) states that the reliability of animal signals can be enforced if those signals cost the signaling creature something in fitness to make (see section 2). The significance of this first experiment to our overall project is that we have made several fundamental improvements to our experimental model. For example, the model of aggression has been revised, in order to bring it closer to current opinion in behavioural biology on the nature and function of aggression in animals (see sections 2 and 3). And we have introduced an evolutionary scenario in which given (i) a large number of possible signaling strategies, and (ii) an initially random distribution of signaling strategies across a population of individuals, the subsequent distribution of signaling strategies evolves in response to ecological context (see section 3).

The second experiment introduces a significant complication into the model. In both experiments, the activity of each animat is, in part, determined by the values of the signals that that animat receives. In the first experiment, an incoming signal is not, in any way, 'interpreted' by the receiver. But it seems likely that, as well as signaling strategies, animals will evolve to have *receiving strategies;* that is, a receiver will give some weight (or degree of importance) to incoming signals. To investigate the effect of receiver tactics on aggressive communication systems, we introduce the concurrent evolution of individual signaling *and* receiving strategies.

2 Biological Background: Aggressive Signaling

Lorenz (1966) characterized aggression as something akin to a 'spontaneous appetite,' such that, in the absence of the performance of aggressive acts, the tendency to behave aggressively increases with time. However, although there are some examples of aggression increasing in some species as a result of social isolation, there is little empirical evidence in favour of the Lorenzian story. Moreover, if aggressive acts did occur purely as a consequence of endogenous changes in the animal, without the correlated presence of any external threat or adaptive goal, then it is hard to see how to make evolutionary-functional sense of aggression. The message is that aggression needs to be

---

[1] There are studies in the A-Life and related literature which, in many ways, foreshadow the SBE-approach. Examples include Koza, Rice, and Roughgarden's (1992) study of foraging behaviour in a population of simulated lizards, and te Boekhorst and Hogeweg's (1994) use of a synthetic ecology to investigate the formation of travel parties in orang-utans.

treated not as an end in itself, but as an adaptive phenomenon, with an adaptive purpose, such as to win or to defend a resource. (For a discussion of all these issues, see (Archer, 1988).)

The phenotypic traits that determine an animal's ability to win a fight are called its resource holding potential ('RHP'). An example of RHP would be size. Signals which are biologically correlated with RHP cannot be faked. For instance, only large toads have low-pitched croaks. A small toad cannot fake a low voice, and, hence, cannot 'lie' about its size with that style of signal (Davies & Halliday, 1978). However, other signals — such as signals of *aggressive intentions* — are not necessarily reliable. Cheats who consistently signaled high levels of aggressive intent, whatever their actual intentions, could well prosper when confronted by 'trusting' opponents. In considering this possibility, Zahavi (1975) has argued that the reliability of intention-signals could be increased if the animal concerned had to invest, in some way, in those signals. This idea — known as *the handicap principle* — is illustrated by the fact that a signal which is, for example, wasteful of energy is, as a consequence of that wastefulness, reliably predictive of the possession of energy; hence honesty is enforced (Grafen, 1990). To be reliable, signals of aggressive intent must be more costly in fitness terms than they strictly need be merely to communicate unambiguously the information at issue. Moreover, the costs involved must be differential. A specific signal indicating a particular level of intended escalation must be proportionally more costly to a weak individual than to a strong individual.

In confrontations involving aggressive signaling (as in many other scenarios), what counts as an adaptively fit individual strategy will be determined by the frequencies with which the various available strategies are adopted by the other members of the population; i.e., individual fitness is frequency-dependent. To investigate such situations, theoretical biologists employing formal models have come to use the concept of an *evolutionarily stable strategy* or 'ESS' (Maynard Smith, 1982). An ESS is a strategy which, when adopted by most members of a population, means that that population cannot be invaded by a rare alternative strategy.

Zahavi's handicap principle has received support from ESS models (e.g., Enquist, 1985; Grafen, 1990). But while ESS theory is, without doubt, a powerful framework for modeling signaling systems, existing ESS models, at least, are limited in their application to natural systems. For example, they do not allow for two-way information flow, or for situations in which there are many receivers of one signal (Grafen & Johnstone, 1993). We shall be concerned with just such a multi-agent context; i.e., a context in which there is two-way information flow and in which a signal is, in general, picked up by many receivers. Hence the experiments presented here are *early steps towards* a situation in which SBE may be of service to biology, by providing a methodology for bridging the explanatory gap between abstract mathematical models and natural environments.[2]

3 Experimental Model

The synthetic world is two-dimensional and non-toroidal; i.e., it is a flat plane, the edges of which are barriers to movement. Its dimensions are 1000 by 1000 units (each animat being round and 12 units in diameter). Space is continuous.

The animats have two *highly* idealized sensory modalities. The 'visual' system is based on a 36-pixel eye providing information in a full 360 degree radius around the animat, with an arbitrarily imposed maximum range of 165 units. Each pixel returns a value corresponding to the proportion of that pixel's receptive field containing other animats. The 'olfactory' system employs principles similar to those used for vision, the only differences being that olfactory range is only 35 units, and that food particles are treated as point sources.

At the start of a run, a number of animats (all with equal energy levels) and a number of food particles (all with equal energy values) are distributed randomly throughout the world (see figure 1). When an animat lands on a food particle, the energy value of the particle is transferred to the animat, thereby incrementing that animat's existing energy level. So, in order to maintain the resource, new food particles are added (with a random distribution) at each time-step. The resource is 'capped' so that the food supply is never more plentiful than at the beginning of the run.

---

[2] For reasons of space, this summary of the relevant biological literature has been highly selective. A more comprehensive review, together with a more complete list of references, appears in (Wheeler & de Bourcier, 1994). That paper also includes a more detailed description of our basic experimental model.

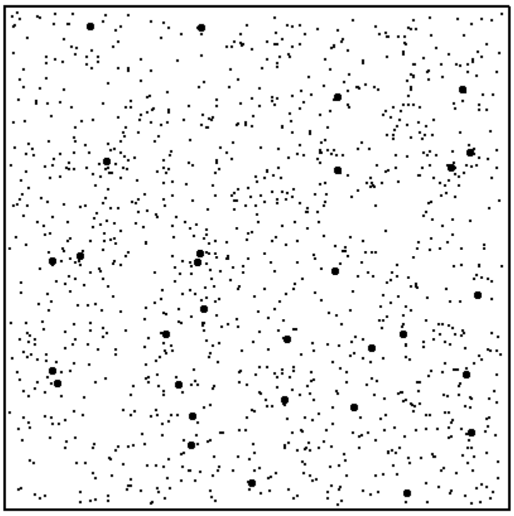

Figure 1: *The world at the start of a typical run. The particles of food are shown as dots, and the animats are shown as filled-in circles. Both food particles and animats are placed randomly.*

Animats pick up up energy from food, but they also pay a series of energy costs (including a small existence-cost deducted at every time step, and costs for fighting, moving, signaling, and reproducing — see below). If an animat's energy level sinks to 0, then it is removed from the world. In consequence, food-finding is an essential task, and, to encourage foraging behaviour, each animat has a hunger level (its disposition to move towards food) which changes in a way inversely proportional to an animat's energy level. Because the food supply is limited, animats are in competition for the available resources.

In addition to a disposition to find food, each animat has a disposition to make aggressive movements (its 'aggression level'). An aggressive movement is defined as one in which the individual in question moves directly towards another animat. If animats touch, a fight is deemed to have taken place, and the combatants suffer large energy reductions. When an animat makes an aggressive movement, its level of aggression is increased by an amount proportional to the previous aggression level. Conversely, a non-aggressive movement results in a decrease in an animat's level of aggression, by an amount proportional to the previous aggression level. Increases in aggression levels take place much more swiftly than decreases (see the parameter-values for the changes in aggression, as detailed in footnote 4).

So the disposition to respond aggressively towards other animats decreases with time in the absence of aggressive encounters, and increases as soon as a confrontation develops (thereby making a further aggressive act more likely in the immediate future). This feature means that we are already working with a non-Lorenzian model of aggression. Moreover, in the ecological situation as a whole, aggression serves an adaptive purpose. Foraging behaviour is essential for survival. So it benefits an individual animat to inhabit an area which is not being foraged by other animats. Aggressive behaviour helps to 'defend' such an exclusive area. Of course, aggressive behaviour is reactively triggered by the presence of other animats within visual range; so animats do not, in any sense, 'plan' their aggression with respect to maintaining a foraging area. The adaptive significance of the behaviour is an emergent property of the ecological situation which we, as external observers, can identify, but which is not explicitly programmed into the animats themselves.[3]

Animats produce aggressive signals whenever at least one other animat is within visual range, and receive the aggressive signals of any other animats within that range. These signals are indicators, to receiving animats, of the *apparent* aggression of signalers (see below).

---

[3] Although we shall not pursue the issue here, we consider the multi-agent dynamics observed in our model to amount to a minimal form of territorial behaviour (see Wheeler & de Bourcier, 1994).

At each time-step, the direction in which each animat will move (one of the 36 directions) is calculated using the following probabilistic equation (each movement results in a small reduction in available energy):

$$p(d) = \frac{h.s(d) + a.v(d) + t(o).v(o) + c}{\sum_{i=1}^{n}(h.s(d_i) + a.v(d_i) + t(o_i).v(o_i) + c)}$$

where *p(d)* is the probability that the particular animat will move in the direction d; n is the number of possible directions of movement; h is the animat's hunger level; *s(d)* is the value returned by the olfactory system in direction d; a is the animat's aggression level; *v(d)* is the value returned by the visual system for direction d; o is the direction 180 degrees off d — i.e., in the opposite direction to d; *t(o)* is the threat (aggressive signal) that the animat perceives from other animats from the opposite direction to d; *v(o)* is the value returned by the visual system in the opposite direction to d; and c is a small constant which prevents zero probabilities.

Roughly speaking, the effects of the movement equation, for an animat, I, can be summarized as follows: (i) the probability of I moving in the direction of food is proportional to I's degree of hunger; (ii) the probability of I moving in the direction of another animat is proportional to I's aggression level; (iii) the probability of I moving away from another animat is proportional to the threat which I perceives from that second animat; and (iv) if there are no other animats in I's visual field, and no food within I's olfactory range, then I will make a random movement.

The various members of the population adopt different *signaling strategies,* which are under evolutionary control. Signals are produced in accordance with the calculation $S = A + ((C/100).A)$ where S is the value of the signal made, A is that individual's current aggression, and C is an individual-specific constant, the value of which is in the range 0-100. For each individual, C is specified by a bit string genotype. At the beginning of a run, a random population of genotypes is created, producing a random distribution of signaling strategies. However, when an animat achieves a pre-defined (high) energy level, it will asexually reproduce. The result is an only child placed randomly in the world. The offspring is given the same initial energy level as each member of the population had at the start of the run, and the corresponding amount of energy is deducted from the parent. The parent's genotype is copied over to the offspring, but there is a small probability that a genetic mutation will take place. So it is possible that the child will adopt a different signaling strategy to its parent. (Throughout the first experiment, an 8 bit genotype was used, and the mutation rate was set to be a 0.05 chance that a bit-flip mutation will occur as each bit is copied.)

Notice that this form of artificial evolution features *no explicit fitness function.* Rather, the selection pressures imposed by the ecological context mean that different strategies will have different fitness consequences. Only those individuals adopting adaptively fit strategies will have a high probability of becoming strong enough, in energy terms, to reproduce.

For the purposes of this investigation, the key selection pressure under the direct control of the experimenter is the cost of signaling. Aggressive signals are displays for which a signaling animat has to pay via a deduction in energy. This cost increases as the level of aggression signaled increases, so that it costs more in energy to make a more aggressive signal. Thus the costs involved are differential, in the sense required by the handicap principle, because, given a specific signal made by a high-energy individual, it will cost a low-energy individual proportionally more to produce the same signal.[4]

4 Results: The Effect of Cost on Signal Reliability

To uncover the trends in signaling behaviour, we partitioned the total population into four sub-populations as follows: The first group included all those individuals producing signals indicating

---

[4] The values of the various parameters were set (largely as a result of trial and error) as follows: initial supply of food = 1200 particles; initial size of population = 30; initial energy level = 300; energy level at which reproduction takes place = 1000; initial level of aggression = 100; energy value of 1 particle of food = 45; rate of food replenishment = a maximum of 16 particles per time step; maximum supply of food at any one time = 1200; existence-cost = 1; movement-cost = 2; cost of fighting = 100 units of energy per time step of fight; increment to aggression level after making an aggressive movement = 1/10 of previous aggression level; reduction in aggression level after a non-aggressive movement = 1/100 of previous aggression level; constant preventing zero probabilities = 1.

Figure 2: *Low Cost Signaling — the total energy levels of the four sub-populations are plotted against time.*

levels of aggression between 0% and 25% in excess of actual aggression, the second 25% to 50% in excess of actual aggression, the third 50% to 75%, and the fourth 75% to 100%. So the 0-25% group produce the most reliable signals, *in the sense that* the signals given out by those individuals are (relative to the signals produced by those in other groups) more accurate reflections of those individuals' dispositions to make aggressive movements. The fourth group are the most extreme bluffers. Each sub-population, at any one time, included all individuals adopting a strategy from the appropriate band, including any offspring.

We ran the simulation many times, setting various values for the cost of signaling. Below we discuss two typical examples of the results obtained.

4.1 Low Cost Signaling

The cost of signaling was set at 0.002 units of energy deducted per unit of aggression signaled. Figure 2 shows plots of the total energy present in each of the 4 sub-populations, during each of the first 16000 time steps. (We take it that 'total energy' is an intuitively satisfactory guide to general adaptive success.)

After an initial settling down period, the 75-100% group tend to dominate. Both the 0-25% group and the 25-50% group struggle to survive in this eco-system, and the 50-75% group mounts only one unsuccessful challenge (see the period around time = 8000). In every one of the many runs of this first experimental set-up, featuring low cost signaling, the most extreme bluffers have been by far the most successful sub-population. So, in this simple eco-system, when there is a low cost to signaling, the most adaptive strategy is to produce signals that indicate levels of aggression well in excess of actual aggression. And, on the basis of empirical observation, extreme bluffing appears to have the general character of an ESS. Following the initial decline of the low-bluffing strategies, these groups fail to re-establish themselves in the population, despite several reappearances due to fortuitous mutations. The results also suggest that reliable signaling would not be an ESS, as a

Figure 3: *High Cost Signaling — the total energy levels of the four sub-populations are plotted against time.*

population of reliable signalers could be invaded by high-bluffing mutants.

4.2 High Cost Signaling

The cost of signaling is now increased to 0.2 units of energy deducted per unit of aggression signaled. Figure 3 shows the results of one representative run. Again the graph shows the total energy present in each of the 4 sub-populations, during the first 16000 time steps.

In this particular high cost run, the 0-25% group and the 25-50% group tend to co-exist alongside each other in the eco-system. For most of the run, the relative positions of the two relevant plots indicate that the 0-25% group is marginally more successful than the 25-50% group.

Over many runs featuring this high cost of signaling, we have observed that the two groups producing the more reliable signals are always dominant. So, in this simple eco-system, when there is a high cost to signaling, it is no longer beneficial to bluff excessively, because the energy cost incurred through such behaviour is prohibitive. However, it must be stressed that it is still adaptive to bluff. The high cost of signaling *restricts* the degree of bluffing, it does not enforce perfectly reliable signaling. But we found that if we increased the cost of signaling much further, in an attempt to enforce increased reliability, the tax on signaling became so harsh, that the whole population quickly died out. We can only speculate that similar situations — where higher costs restrict, but do not eliminate cheating — may occur in some natural ecologies.

Because the two groups in question cover such a wide spread of signaling strategies, the ESS-concept can be used in only a vague, intuitive sense. However, it is clear that once the two higher-bluff strategies have died out, they are not able to re-establish themselves against populations made up of more reliable signalers, despite several reappearances through mutations.

So, in this simple synthetic ecology, a high cost to signaling does result in *increased* reliability of signals. This suggests that the *general* logic of the handicap principle can carry over to the sort of multi-agent signaling systems not covered by existing ESS models.

5 The Introduction of Receiving Strategies

The conclusions from our first experiment constitute the beginning, rather than the end, of a story. To model more realistic signaling systems, we need to consider not only the behaviour of signalers, but also the behaviour of *receivers*. 'Receiver psychology' has become an increasingly important issue in the biological literature (e.g., Stamp Dawkins & Guilford, 1991; McGregor, 1993).

In the experiment described above, the threat values for the movement equation were simply the values of the incoming aggressive signals. However, we now extend the experimental model, so that each animat has not only a genetically specified signaling strategy, but a similarly specified receiving strategy. Each animat's receiving strategy is determined by an individual-specific constant, K that 'weights' incoming signals, in order to produce the threat value for the movement equation. So $T = R.(K/100)$, where T is the threat value in the movement equation, R is the incoming signal, and K is an individual-specific constant, in the range 0-200. A K of 0 would result in that individual ignoring incoming signals; a K of 100 means that the value of the incoming signal itself is used as the threat value (i.e., equivalent to the situation in the first experiment); and a K of 200 results in incoming signals being doubled, and that resulting value being used as the threat value. Henceforth we shall speak of this weighting system as the *degree of importance* assigned to incoming signals (the higher the value of K, the higher the degree of importance). So when an individual signals, the exact effect that that signal has on the behaviour patterns of animats within visual range (what that signal 'means' to receivers) is now also a function of those animats' receiving strategies.

To maintain a balance between the range of possible receiving strategies, and the range of possible signaling strategies, we now allow the existence of individuals who produce signals indicating levels of aggression *lower* than actual aggression. Individuals produce signals according to the calculation $S = A.(C/100)$, where S is the value of the signal made, A is that individual's current aggression, and C is an individual-specific constant, in the range 0-200. A C of 0 is equivalent to not making any signal, a C of 100 is equivalent to producing indicators of actual aggression, and a C of 200 is equivalent to producing signals indicating twice actual aggression.

The length of each individual's genotype is doubled, and encodes for both a signaling and a receiving strategy. We employ the same process of artificial evolution (with the same mutation rate) as was used in the first experiment.

6 Results: Interacting Strategies

Once again, the population was divided into four sub-populations on the basis of signaling strategy. The groups were identified by ranges in the value of the individual-specific signaling-constant, C (Group 1: 0-50, Group 2: 50-100, Group 3: 100-150, Group 4: 150-200). The total energy present in each sub-population was then recorded against time. However, for this experiment, we also divided the population into four sub-populations on the basis of receiving strategy, and recorded the total energy present in each of those groups as well. The receiving groups were identified by ranges in the value of the individual-specific receiving-constant, K (Group 1: 0-50, Group 2: 50-100, Group 3: 100-150, Group 4: 150-200).

Our work on the interaction between signaling and receiving strategies is still in its early stages, and we do not pretend to have anything approaching a full understanding of the relatively complex population dynamics that unfold at different costs of signaling. Much more analysis needs to be done. However, there are some clear trends in the observed behaviour, and below we discuss the results of two suggestive runs.

6.1 Low Cost Signaling

The cost of signaling was set to be low (as before, 0.002 units of energy deducted per unit of aggression signaled). Figure 4 shows plots of the total energy present in each of the 4 receiving sub-populations, over the first 64000 times steps of the run. Figure 5 shows the equivalent plots for each of the 4 signaling sub-populations. In the following discussion, we shall often speak of 'signalers' and 'receivers'; but it should be remembered that each individual is both a signaler and

Figure 4: *Low Cost Signaling — the total energy levels of the four* **RECEIVING** *sub-populations are plotted against time. See figure 5 for plots of the concurrently evolved signaling strategies.*

a receiver. (Notice that the time period over which these graphs were plotted is four times that used in the first experiment. Given the increased complexity of the scenario, it takes much longer to identify significant trends in the population dynamics.)

The dominant signaling strategy for most of the period shown (and for low cost signaling cases in general) is signaling group 4 — the high bluffers (see figure 5). This was expected, given the low cost of signaling. However, it might have been thought that the failure of signals to reflect actual aggression would mean that it would pay individuals to ignore those signals, so that receiving group 1 would be the most successful. Then, with signals having little (if any) effect on receivers, the communication system would be in danger of breaking down (cf. Krebs & Davies, 1987). But, in fact, the dominant receiving strategy for most of the run is receiving group 4 — the group that gives the *highest* degrees of importance to incoming signals (see figure 4). The explanation for this seems to be that there is, in this simple ecology, just one source of information about the quality of potential adversaries, namely aggressive signals. So individuals do not have the option of switching to another, more reliable source of information, such as RHP. Under these circumstances, the strategy of giving a high degree of importance to incoming signals is less likely to lead a receiver into potentially costly fights. (The extent to which an individual ignores incoming signals is the extent to which that individual's behaviour towards other animats will tend to be driven directly by its own aggression level.) In that sense, there is selective pressure on receivers to give a high degree of importance to incoming signals. However, there is another side to the adaptive story.

There are periods when the dominant signaling and receiving strategies are challenged by signaling group 3 and receiving group 3. This is common in low cost runs, and can be seen most clearly in figures 4 and 5 in the period following t = 46000 (although there are less dramatic examples throughout). Receiving group 3 temporarily displaces group 4 as the dominant receiving strategy. We believe (from watching the distribution of strategies change whilst observing the aggressive interactions during such periods) that the explanation for this is as follows: Individuals who evolve to give less weight to incoming signals make a larger number of aggressive movements

Figure 5: *Low Cost Signaling — the total energy levels of the four* **SIGNALING** *sub-populations are plotted against time. See figure 4 for plots of the concurrently evolved receiving strategies.*

(because incoming signals produce lower threat values). This, in turn, increases the aggression levels of the animats adopting that policy (because when an animat makes an aggressive movement, its level of aggression is increased); so those animats produce bigger signals. Individuals giving high degrees of importance to the incoming signals then tend to retreat more often, giving the advantage (in terms of foraging areas) to those individuals adopting receiving strategies in group 3. There is then an improvement in the fortunes of signaling group 3. Why should this happen? On this point, we can, at present, offer no more than a conjecture. But it is plausible that it occurs as a direct result of the increase in the general aggression levels of individuals in receiving group 3. Higher aggression levels mean bigger signals; so a lower value of C (the individual-specific signaling-constant) will now produce signals equivalent to those previously produced with a higher value of C. Individuals with a lower C obtain a small saving in energy (and, thereby, a small fitness advantage) by signaling at a lower level. However, once the majority of the population have adopted the newly dominant signaling and receiving strategies, the adaptive advantage is lost, and the new arrangement is unstable against invasion by mutants who adopt the signaling and receiving strategies which are generally dominant in low cost ecologies (i.e., signaling group 4/receiving group 4).

One can think of the two strategic arrangements — (i) signaling group 4/receiving group 4 and (ii) signaling group 3/receiving group 3 — as attractors in the dynamics of the self-organizing ecological system, with the first the much stronger of the two. The probabilistic behaviour, together with the ongoing fluctuations in the levels of aggression (and, therefore, the signals being produced), mean that the system occasionally evolves towards the weaker attractor.

6.2 High Cost Signaling

As in the first experiment, the cost of signaling was then increased to 0.2 units of energy deducted per unit of aggression signaled. Figure 6 shows plots of the total energy present in each of the 4

Figure 6: *High Cost Signaling — the total energy levels of the four* **RECEIVING** *sub-populations are plotted against time. See figure 7 for plots of the concurrently evolved signaling strategies.*

receiving sub-populations, over the first 64000 times steps of the run. Figure 7 shows the equivalent plots for each of the 4 signaling sub-populations.

With a high cost to signaling, there is a powerful selection pressure for individuals to minimize the amount of energy that they expend on signaling. So it might seem relatively unsurprising that the dominant signaling group for most of the run is group 1 — the group producing the lowest-value signals (see figure 7). But now notice that these individuals are producing signals indicating levels of aggression *much less than* actual aggression. From this one might be led to conclude that the introduction of receiving strategies upsets the operation of the handicap principle. Signals seem to be unreliable. Moreover, because signals are indicating apparent levels of aggression which are, in fact, less than actual aggression, there should be a higher probability that individuals will be drawn into costly fights. But, in fact, the scenario does *not* lead to an escalation in the number of fights; and, once we take account of the interaction between the signaling and receiving strategies, it becomes clear that something very like the handicap principle *does* operate.

The dominant receiving group for most of the run is group 4 — the group giving the *highest* degrees of importance to incoming signals (see figure 6). This is significant, because there is a theoretical equivalence, as far as the reactions of receivers are concerned, between situations where signals are small compared to actual aggression, but receivers give high degrees of importance to those signals, and situations where signals are higher, but receivers give those signals less importance. As an over-simplified example, consider that a one-way dyadic interaction in which the signaler's $C = 50$ and the receiver's $K = 200$ is equivalent to a similar interaction in which the signaler's $C = 100$ and the receiver's $K = 100$. Both result in a threat value for the receiving individual equal to the actual aggression of the signaling individual. But the key difference is that the signaling animat in the latter case would expend much more energy through signaling. Of course, there are ranges of different signaling and receiving strategies operating in the population at any one time; and communication events are two-way and, in general, multi-agent. Hence the situation is never that straightforward. But the point is that the ecology tends to evolve towards a

Figure 7: *High Cost Signaling — the total energy levels of the four* **SIGNALING** *sub-populations are plotted against time. See figure 6 for plots of the concurrently evolved receiving strategies.*

situation which, in terms of receiver-behaviour, is *equivalent to* reliable signaling. And the adaptive advantage of that evolutionary solution is that the costs incurred by signaling animats are much lower than they would be if signals were reflections of actual aggression.

But now we should qualify that last statement, by pointing out that the communication system as a whole is not entirely free of 'cheating'. For example, groups of signalers from the three higher-signaling groups do enjoy some adaptive success. The reason for this is that for most of the run, receivers are giving a high degree of importance to incoming signals. So, even where costs are high, there may well be some occasional advantage to producing bigger signals, which will have a high pay-off in terms of the subsequent behaviour of receivers.

7 Conclusions

We hesitate to draw bold conclusions about the natural world on the basis of our simple synthetic ecology. However, the evidence lends support to the hypothesis that where the fitness costs of signaling are low, signalers will tend to produce signals indicating levels of aggression well in excess of actual aggression. *If there is no alternative source of relevant information,* receivers may still pay heed to those signals. Where the fitness costs of signaling are high, the pressure on signalers to reduce the level of signaling may still lead to communication systems in which signals are not direct reflections of quality, in that signalers may tend to produce signals indicating levels of aggression *lower* than actual aggression. However, if receivers evolve to give a high degree of importance to those signals, the effect would be to *compensate* for the actual values of the signals. As receivers, individuals would still behave *just as if* signals were direct reflections of aggression (so they would benefit from not being drawn into costly conflicts); as signalers, individuals would benefit from the low level of signaling.

Our work on the interaction between signaling and receiving strategies has only just begun. Even in this simple synthetic ecology, the dynamics of the communication system are made relatively complex by the need for animats to forage, the probabilistic nature of the behaviour, the constantly changing aggression levels, and the fact that many different signaling and receiving strategies can coexist in the population at any one time. We believe that these features of the experimental model are crucial in ensuring that the results of the simulation are non-trivial. But these same features complicate the process of explanation. Thus, whilst we believe that our proposed explanations of the observed behaviour are fundamentally correct, our conclusions would be strengthened by further analysis. That is our next task. In particular, we need a more detailed understanding of the ways in which the system is sensitive to changes in the values of parameters other than the cost of signaling.

Aggressive communication is an adaptive phenomenon about which there are unanswered questions. We hope that our ongoing work in SBE will help to find some of those answers.

References

Archer, J. (1988). *The Behavioural Biology of Aggression.* Cambridge University Press, Cambridge.

Davies, N. B., & Halliday, T. R. (1978). Deep croaks and fighting assessment in toads *Bufo bufo. Nature*, *274*, 683–5.

de Bourcier, P., & Wheeler, M. (1994). Signalling and territorial aggression: An investigation by means of synthetic behavioural ecology. In Cliff, D., Husbands, P., Meyer, J.-A., & Wilson, S. W. (Eds.), *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, pp. 463–72 Cambridge, Massachusetts. M.I.T. Press / Bradford Books.

Enquist, M. (1985). Communication during aggressive interactions with particular reference to variation in choice of behaviour. *Animal Behaviour*, *33*, 1152–1161.

Grafen, A. (1990). Biological signals as handicaps. *Journal of Theoretical Biology*, *144*, 517–546.

Grafen, A., & Johnstone, R. (1993). Why we need ESS signalling theory. *Philosophical Transactions of the Royal Society: Biological Sciences*, *340*, 245–250.

Koza, J. R., Rice, J. P., & Roughgarden, J. (1992). Evolution of food-foraging strategies for the caribbean *anolis* lizard using genetic programming. *Adaptive Behavior*, *1*, 171–200.

Krebs, J. R., & Davies, N. B. (1987). *An Introduction to Behavioural Ecology* (2nd edition). Blackwell Scientific, Oxford.

Lorenz, K. (1966). *On Aggression.* Methuen, London.

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

McGregor, P. (1993). Signalling in territorial systems: a context for individual identification, ranging and eavesdropping. *Philosophical Transactions of the Royal Society: Biological Sciences*, *340*, 237–244.

Stamp Dawkins, M., & Guilford, T. (1991). The corruption of honest signalling. *Animal Behaviour*, *41*(5), 865–73.

te Boekhorst, I., & Hogeweg, P. (1994). Effects of tree size on travelband formation in orang-utans: Data analysis suggested by a model study. In Brooks, R., & Maes, P. (Eds.), *Proceedings of Artificial Life IV*, pp. 119–129 Cambridge, Massachusetts. M.I.T. Press.

Wheeler, M., & de Bourcier, P. (1994). How not to murder your neighbor: Using synthetic behavioral ecology to study aggressive signaling. Tech. rep. 357, School of Cognitive and Computing Sciences, University of Sussex. Submitted to the journal *Adaptive Behavior.*

Zahavi, A. (1975). Mate selection — a selection for a handicap. *Journal of Theoretical Biology*, *53*, 205–214.