In 4th European Conference on Artificial Life 97,Brighton UK

# CHORUSING AND CONTROLLED CLUSTERING FOR MINIMAL MOBILE AGENTS

Owen Holland, Chris Melhuish, and Steve Hoddell
Intelligent Autonomous Systems Engineering Laboratory, Faculty of Engineering,
University of the West of England, Coldharbour Lane, Bristol BS16 1QY, UK
email: o-hollan/cr-melhu@uwe.ac.uk

## ABSTRACT

This paper examines some strategies for controlling the sizes of clusters of simple mobile agents. The basic method is loosely modelled on the signalling behaviour of species such as crickets, frogs, and fireflies, where males attract females by making periodic calls or light flashes. A characteristic of such behaviour is that groups of animals broadcast their signals in synchrony, forming a chorus. In this paper, synchrony is used in conjunction with random deviations from synchrony in order to enable each individual to estimate the size of the group over a period of time. By arranging for an individual to approach a group which is below some required size, and leave a group which is above that size, cluster size can be controlled. Two types of controlled clustering are examined: seeded clustering, where a group is required to form at a particular spot, and unseeded clustering, where there is no preferred site. Results from simulations are presented. For seeded clustering, they show reasonable performance for low levels of noise, breaking down as noise levels are increased; for unseeded clustering, effective control is only achieved with almost no noise, and with a heavily modified algorithm.

## 1 Introduction

Multi-agent and multi-robot systems can often be divided into two broad categories: those using symbolic representations and explicit communication; and those using behaviour based techniques and implicit communication 'through the world'. For those developing the second type of system, there is usually some attraction towards minimalism, and a tendency to explore the use of behavioural strategies known to be used by biological systems such as social insects. The system described in this paper is both minimalist, and inspired by the collective behaviour of biological systems. A discussion of the relationships between minimalism, biological inspiration, collective behaviour, and the construction and application of multi-robot systems may be found in (Holland & Melhuish [1997].

### 1.2 Aggregation and the control of group size

Aggregation was identified early on as one of the primitives of swarm systems and collective behaviour (Mataric [1992]). We distinguish two extreme forms: pseudoswarming, in which each individual moves towards a given point using information which is independent of the locations of other agents; and true swarming, in which an individual's movements are wholly determined by the locations of other agents. (Holland & Melhuish [1996b]). Combinations of these types are frequent: leaf cutter ants are at first attracted by the vibrations made by an ant which has found a suitable leaf, but once close enough to detect the leaf are attracted by the leaf itself; robots attracted to an infra-red source can change course to avoid one another.

The control of group size has received little attention. It may be required for a number of reasons. For example, if a localised resource is sufficient for only a limited number of agents, there is little point in attracting extra agents to the resource. Again, it has been established by at least two sets of robot experiments (Beckers et al [1994]; Fontan & Mataric [1996]) that there may be an optimum number of robots for carrying out a given task under given circumstances; in such cases, the control of the size of the group undertaking the task may be critical to achieving the task quickly or efficiently.

The aim of this work is to find a suitable method for regulating the size of a group of agents. The basic idea is that each agent should individually derive an estimate for the size of any group of which it is effectively a member, and should then either approach the centre or focus of the group if the estimated group size is less than or equal to some internal parameter expressing the 'desired' group size, or should move away from the centre or focus if the estimated group size is too large. If the group is required to form at a particular point, then we assume that at that point sits a beacon which emits some omnidirectional field with a strength which reduces with distance from the beacon. Motion away from the group can therefore be arranged by moving away from the beacon. We call this seeded clustering. (The agents do not need to be able to sense the orientation of the beacon directly; the method of klinokinesis can be used - see Holland & Melhuish [1996a]).Where there is no requirement to cluster in a particular location, no beacon is available as a guide for movement, and so some signal must be derived from sensing or picking up signals from the other agents in the cluster. We call this unseeded clustering, and it is a much more difficult and complex problem.

A very simple method of achieving seeded clustering is to arrange for the agents to reduce the intensity of the attractive signal by the merely passive fact of their presence. This was used by Kube and Zhang [1992] in an early study of aggregation; robots were attracted to a lighted box in an arena, but robots reaching the box blocked the light from the view of the other robots, resulting in no more robots being attracted to the box once there were enough there to completely block the light. This use of passive properties is an attractive and neat solution, but works only when the number of agents required just happens to be the number obtained; it is difficult to tune, and is a rather precise function of the nature of the environment, task, and agents, so it cannot be extended to serve as a general method. Active properties offer more potential generality, and we have adopted the following set of constraints: the agents will be able to transmit and receive some actively broadcast signal, omnidirectionally transmitted and detected, and decreasing in intensity with some function of distance; and each agent will be assumed to contain some internal parameter which indicates the size of group required.

Some strategies satisfying this restriction can be ruled out after some quite general considerations. One factor of interest is the power of the signal. (For an insightful study of many aspects of signalling in animals, including power, see Endler [1993]). Since the method of regulation is constrained to be active, each broadcast by a agent will use power, so in the interests of economy some means of reducing the power would be useful. Further, the range will be a function of the broadcast power; a long range would require a high power handling capacity, which is likely to use more structural resources than a lower power arrangement. An intermittent broadcast at high power will achieve range, at the cost of losing temporal granularity; this intermittent high power need not require the capacity for high power generation, as some accumulator mechanism can be used to store energy which is suddenly released to give a high instantaneous power, just as the flea winds its legs up over a period of time and releases them to make a leap. We therefore decided to explore the use of an intermittent signal, and ruled out looking at the summation of constantly transmitted signals, for example from agents within a certain distance of the source.

The simplest intermittent signal is the equivalent of a click. It is characteristic of a click that two clicks will never overlap. If we assume that the intensity of a click must exceed some threshold in order to be registered by an agent, this means that the range over which a click can be detected will be fixed at some maximum. In contrast, a longer signal which overlaps other such signals can be expected to summate, and so will typically be detectable at a greater distance than a single signal, other things being equal. Although it would have been quite simple to devise a clicking mechanism - for example, clicking when close to the source, leakily integrating the time series of clicks, and moving away from the source if the integrated quantity was higher than the group size threshold - this would have a fixed maximum range of action from the source. It is desirable that the message that the group is already large enough (or too large) should be detectable at the largest possible range, and so it was decided to investigate signals which were brief (saving power) but which could gain in range by being superimposed. This train of thought led us to study reports of natural systems which used such brief repeated signals. There are many such systems in nature; the best known are probably the sound choruses of crickets and frogs, and the light flashes of fireflies.

### 1.3 Chorusing in crickets, fireflies, and frogs

There is a large and fascinating literature on chorusing in crickets, fireflies, and frogs; particularly useful and accessible texts are: Greenfield [1994]; Sebeok [1977]; Ryan et al [1981]; Ewing [1989]; and Alexander [1975]. The outstanding characteristic for us was that most of the creatures which chorus, or broadcast intermittent signals in synchrony, appear to use a similar mechanism for synchronisation. Greenfield calls this mechanism ‘the basic phase delay interactive algorithm’; essentially, a sawtooth pacemaker which produces a chirp when it rises to its maximum level is reset to the basal level by the perception of a chirp from another animal. If both animals have the same pacemaker period, they will be in synchrony on the subsequent chirp of the interrupting animal. There are of course many variations on this theme, including alternating rather than synchronous chirping, but most chorusing appears to follow this rule. The function of chorusing can be very varied, and is not always clear. One clear function is that of increasing the range of a group signal while retaining the temporal features which enable it to be identified and discriminated from other signals using the same modality. The role of the signal itself may be to attract females for mating, or to confuse predators by making signal localisation difficult. However, we have found no mention of chorusing being used to regulate group size.

## 2 Methods for estimating the size of an agent cluster

Our first idea for estimating group size was to take the local summed intensity of calls as the representation of group size. Unfortunately, assuming an inverse square law for the fall off of intensity with distance, this varied too strongly with distance to be useful. The second idea was slightly more contrived, but seemed to take advantage of the intermittent structure of the calls. We reasoned that, if all the members of a group were in perfect synchrony at the onset of one chirp, then in any real system there might be some stochastic variation in the subsequent individual start times for the next chirp. If this were ‘corrected’ by the basic phase delay interactive algorithm, then the next chirp would again be subject to some variation in individual start times, and so on. If all individuals had the same distribution of chirp periodicity (modelled by a fixed refractory period plus a random interval) then each individual in a group of *n* could be expected to be the first to chirp on a proportion (1/*n*) of occasions. If a group size of *p* was required, it would only be necessary for each agent to have an internal representation of (1/*p*); if it detected that it was the first to chirp on a proportion of occasions less than (1/*p*), the group must have more than *p* members, and so the agent should select the behaviour of heading away from the group. Of course, the resolution required to control large group sizes would probably be prohibitive, but this looked a promising enough mechanism for controlling small groups.

The proposed mechanism depends on bringing neighbouring agents into synchrony. Since our system required agents to make correct decisions rapidly, it was not thought appropriate to use the basic phase delay interactive algorithm, because this never allows a chirp to be brought forward, but instead works by delaying chirps. Instead, we proposed that an agent able to chirp (i.e. not in a refractory state) would always be induced to chirp immediately it detected another agent's chirp. If one of a group of agents were to chirp, then all other agents not in refractory states would chirp at the next instant; this would provide the fastest possible synchronisation.

## 3 The λ mechanism: simulation details

The simulations use a circular arena 1200 units in diameter, with a source of attraction at the centre. Each agent has two directionally sensitive 'eyes', one on each side of, and at 60° to, the agent's longitudinal axis; the eyes each have a 120° field of view, and the intensity of the source of attraction detected by each eye falls off as the inverse square of the distance from the source to the eye. The environment and the eyes themselves also contribute Gaussian noise to the eyes' readings. Each agent also has a chirp transmitter and receiver; both are omnidirectional, and the detected intensity of a chirp declines with the inverse square of distance. There is a fixed threshold for the detection of any chirp energy.

In the first simulations, it became clear that the algorithm worked to some extent, but that certain modifications would be beneficial. The first change was to limit the chirping to agents within a certain distance of the source; agents outside that distance would synchronise their internal timers to received chirps, but would not themselves emit chirps. We called this mode of operation 'silent mode'. If this was not done, then randomly formed groups of agents at a distance from the source appeared to be able to disrupt the behaviour of the group close to the source.

Another factor noticed in early simulations was the effect of adding some background noise to the reception of chirps. It can act in two ways: it can produce what has become known as stochastic resonance, occasionally enhancing subthreshold chirp signals so that they are registered, and so adding information to behaviour; on the other hand, large amounts of noise can reach the threshold, introducing spurious chirps and so adding noise to behaviour. Any systematic study should therefore include such noise to allow these factors to operate.

The agent can exhibit two related types of movement in relation to the source, each expressed in simple rules. In both, the first step is to compare the sensed source intensities in the left (L) and right (R) eyes. In the first, the attractive behaviour, the rule is:

**If L > R rotate left 60° and move 4 units**
**else rotate right 60° and move 4 units**

This deliberately crude and noisy taxis rapidly brings an agent close to the source and keeps it there. In the second, repulsive behaviour, the rule is:

**If L > R rotate right 60° and move 4 units**
**else rotate left 60° and move 4 units**

This will rapidly move an agent away from the source towards the periphery.

At any time, the rule to use is determined by the group size estimator, or λ mechanism. λ is the estimate of the reciprocal of the group size (1/*n* for a group size *n*). It is compared with τ(λ), the reciprocal of the target group size. If λ < τ(λ), the repulsive rule is used; otherwise, the attractive rule is used.

The parameters of the chirping cycle and λ mechanism are as follows. Each agent has a refractory period (R) of 20 time ticks, during which it cannot chirp. It then passes into the latent period (L). L lasts a maximum of 20 time ticks, and, if no chirp is previously detected from another agent, the agent will start to chirp on one of these time ticks preselected randomly with equal probability. If a chirp from another agent is detected before the time for the spontaneous chirp, the agent will chirp on the next time tick. The chirp period (C) lasts for 6 time ticks. A chirp is transmitted only when the agent is within a radius of 50 units of the source. λ is estimated as the ratio of the number of unstimulated chirps produced to the total number of chirps produced in the previous 500 time ticks.

The simulations presented here examine the performance of the algorithm in an environment containing 15 agents which are initially distributed randomly throughout the arena. Five target group sizes are used (2, 4, 6, 8, and 10) and two conditions of background noise for the chirp mechanism (0% and 3% of maximum intensity). Ten trials were run under each condition. A trial is scored by noting the number of agents within 50 units of the source after 80,000 time ticks. The results are shown in Figure 1.

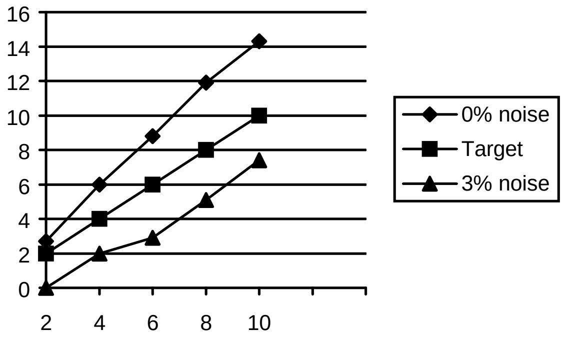

**Figure 1: mean terminal group size**

Qualitative observations of the agents under all conditions were unsurprising. On start up, those agents closest to the source reach it first, form the target group size, and then form a group larger than the target. The group expands a little until some agents are lost, and then a state of flux ensues, with agents being exchanged between the group close to the source, and a highly mobile cloud of agents slightly further away. While sometimes a definite annulus of agents can be seen surrounding the central group, it is not generally visible.

One clear problem is that the agents far away from the source are moving almost randomly, because the low intensity of the source at that distance is swamped by the Gaussian noise in the environment and in the eyes, and so the agents cannot move reliably either towards or away from the source, regardless of the rule they are supposed to be obeying..

As can be seen from Figure 1, the 0% noise condition produces groups rather larger than the target, and the 3% noise produces groups rather smaller than the target. Two conclusions can be drawn from this: setting $\tau(\lambda) = (1/n)$ does not produce accurate group sizes even with no noise, and so $\tau(\lambda)$ should be determined empirically, unless a better mathematical model is formulated; and $\tau(\lambda)$ must be adjusted to take account of any background noise which might affect chirp reception. Nevertheless, the mechanism has been shown to be capable of regulating group size in a systematic way, and so is probably worth investigating further.

These early indications of promise led to an examination of ways in which the performance could be improved while remaining within the limits imposed by the philosophy of minimalism. The first change was to replace the unweighted moving-average computation of $\tau(\lambda)$ with a time-weighted average, which would allow more recent data to be weighted more heavily. For convenience, we changed the factor of interest to the count of the number of first chirps by other agents heard between successive first chirps of the agent in question. On average, the count will be one less than the number of agents in the group, and so a suitable estimate of group size ($\gamma$) is (chirp count + 1). In order to produce a time weighted average, the new estimate of this statistic, $\gamma_t$, is formed from the previous estimate $\gamma_{t-1}$ by setting $\gamma_t = \boldsymbol{a}\gamma + (1-\boldsymbol{a})\,\gamma_{t-1}$. (For $\boldsymbol{a}=1$, this of course reduces to using the estimate $\gamma$, which is actually a simplification when compared to the $\lambda$ mechanism.)

The logic of using a chorusing strategy had been rooted in the effects of summation in extending the distance at which a (combined) signal could be sensed and distinguished from noise. To check that this actually delivered some benefit, we also modified the $\lambda$ mechanism so that summation did not increase range - that each agent transmitted the analogue of a click over a fixed range. We called this the $\kappa$ mechanism.

## 4 Seeded clustering - comparing λ, κ, and γ

In early trials with both the $\lambda$ and the $\gamma$ mechanisms, we noticed that 'full' clusters which appeared stable could be completely destabilised by the approach of a single extra agent. Close examination revealed that this was due to the fact that the arrival of such an agent would cause most or all of the agents in the cluster to begin moving away from the beacon, because they had registered that the cluster was too large. This widening of the cluster would continue until one or more agents (usually several) had moved out of range of the others; the cluster would then begin to collapse again, possibly still including the intruder, and the process of expansion and contraction might continue for several cycles. It would obviously be desirable to make it possible for the intruder or intruders to be repelled or expelled, in preference to individuals already in the cluster.

Two methods were tried. The first was derived from the well-known strategy of orthokinesis, used by many single celled organisms to preferentially exploit regions of higher food concentration. If the speed of movement of an organism moving at random is a decreasing function of food concentration, then the time the organism will spend in a given region will increase with the region's food concentration. In our version, the agent step length was made a function of the sensed beacon strength (which decreases with distance); this means that agents on the outer edges of a cluster are much more volatile than those close to the centre. It also produces much more tightly grouped clusters, which in turn increases the distance between an intruder and the cluster at the time when the intruder's effects are felt. The second method was a silent mode: any agent within a radius of 25 units of the beacon was silent if at the previous step it had calculated that it was in a cluster which was too large. Such an agent should be leaving the cluster, and, by not chirping, it should prevent other cluster members from also deciding that they should leave the cluster.

## 5 Results for seeded clustering using λ, κ, and γ

The first set of experiments was designed to compare the basic λ, γ, and κ mechanisms, for a range of cluster demand sizes, and in both noise-free and noisy environments. Silent modes and orthokinesis were used. The variation in step length for orthokinesis was calculated as follows:

$$\text{step length } \sigma = \max\sigma\left\{1-\left(\frac{\max[LorR]}{\max input}\right)\right\}$$

where $L$, $R$ are the inputs to the left and right sensors, and maxinput is the maximum possible input to either sensor (the input at the beacon site).

The factors of interest were the number of agents clustered around the beacon, and the mean distance from the beacon of those agents outside the cluster. Since stable clusters, when formed, would always fit within a circle of radius 25 units around the beacon, the cluster size was defined as the number within that radius, and any agents outside that radius were defined as being outside the cluster. (In the simulations described in section 3, the cluster radius was taken as 50 units; the reduction to 25 improves spatial resolution, and makes it less likely that an agent leaving the cluster is accidentally counted as a cluster member.) This simplified what might otherwise have been a formidable task of analysis. The (Gaussian) noise levels at the agents' beacon sensors were given mean levels of 0%, 3%, and 6% of the maximum beacon strength. Ten trials were made for each combination of noise level, mechanism type, and cluster demand size. The results are shown in Figs 2 - 4.

One possible problem with using the same fixed population size throughout is that we cannot be sure we are studying the effects of (target) cluster sizes, because the (target) number of non-cluster agents also varies systematically. However, since we are using silent modes, it is generally true that non-cluster agents do not affect cluster agents, and so we would not expect the number of non-cluster agents to affect our recorded outcomes to any degree. Additional simulations have confirmed this.

Figure 2a shows that, for no noise, all three mechanisms are able to control cluster size effectively, with the κ mechanism coming closest to the nominal cluster size. However, Figure 2b reveals the advantage of using the chorusing mechanism, as both the λ and γ mechanisms are able to separate the non-cluster agents from the cluster by up to 168 units, whereas the κ mechanism achieves only around 40 units.

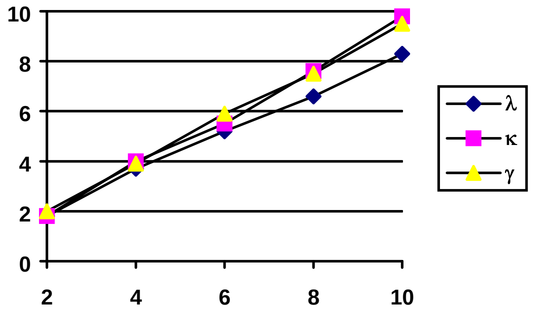

Fig 2a: Mean cluster size - 0% noise

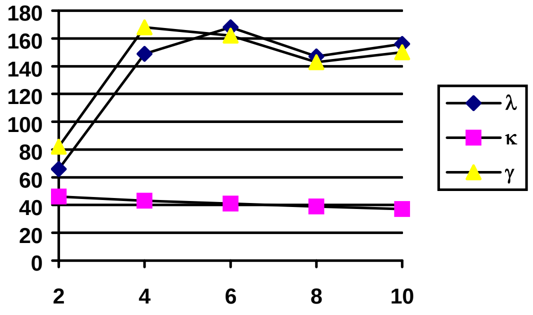

Fig 2b: Mean distance of non-cluster agents - 0% noise

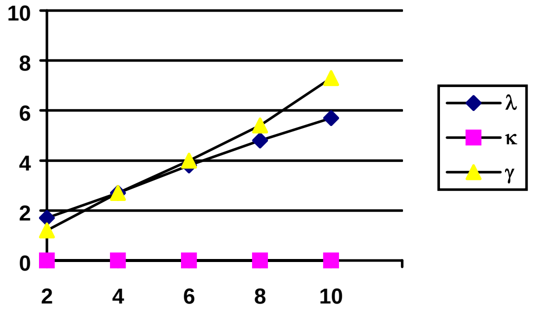

Fig 3a: Mean cluster size - 3% noise

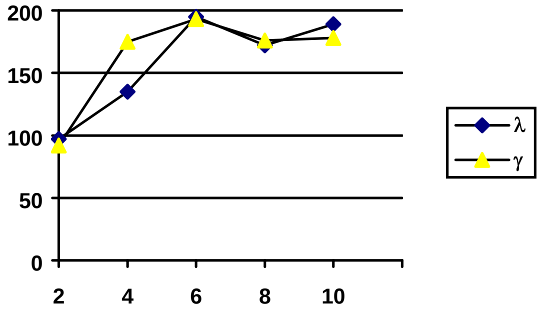

Fig 3b: Mean distance of non-cluster agents - 3% noise

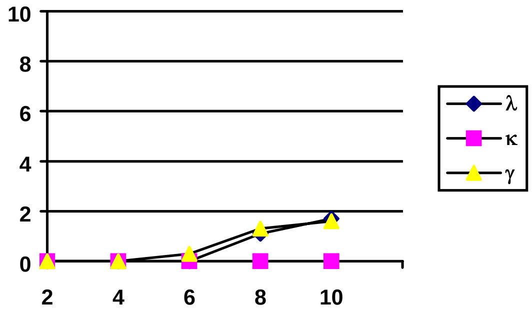

Fig 4a: Mean cluster size - 6% noise

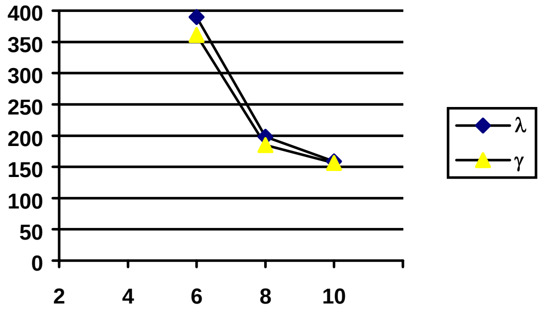

Fig 4b: Mean distance of non-cluster agents - 6% noise

Figure 3a shows that even 3% of noise causes the κ mechanism to fail completely. The λ and γ mechanisms are still able to produce cluster sizes which, although lower than the nominal demand size, are still well differentiated from the higher and lower cluster sizes. Figure 3b shows that both mechanisms achieve large and similar separations between non-cluster and cluster agents - up to 195 units. Figure 4a shows the effective breakdown of all mechanisms, with both the λ and γ mechanisms producing clusters well below the demand size, with the non-cluster agents separated from the cluster agents by very large distances. On the basis of these results, we decided that further development would be confined to the γ mechanism alone.

## 6 Effects of orthokinesis, silent modes, and time-weighted filters on γ mechanism

A further set of experiments examined the effects of orthokinesis, silent modes, and time-weighted filters on the γ mechanism. All parameters were as for the previous set of experiments. The variations studied were: no orthokinesis, with two fixed step sizes of 1 unit (FL1) and 12 units (FL12); orthokinesis with no time-weighted filter, or ***a***=0 (ORT-F); and orthokinesis with no silent mode (ORT-S). Results for these are shown in Figs 5 - 7 along with results for orthokinesis with filter and silent mode (ORT).

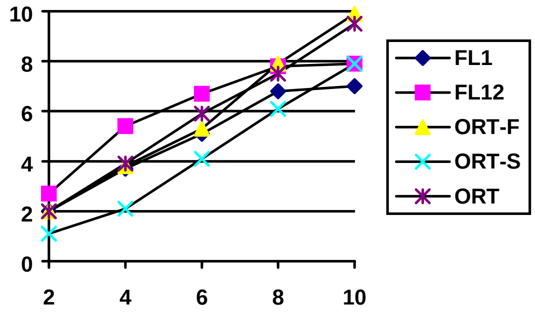

Fig 5: Mean cluster size - 0% noise

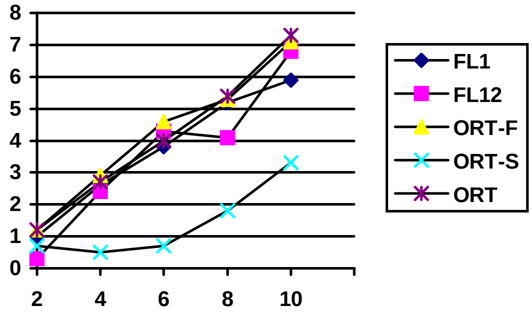

Fig 6: Mean cluster size - 3% noise

From Fig 5, it appears that the strategies which do not use orthokinesis (FL1 and FL12) may run out of resolution for the higher demand cluster sizes. All of the orthokinesis strategies perform adequately, with ORT and ORT-F producing cluster sizes closest to the nominal size.

Fig 6 shows the effects of 3% noise on performance. ORT-S is badly affected; ORT and ORT-F are somewhat degraded, but still have plenty of resolution. This confirmed that the silent mode is important for coping with noise, and that the time weighting introduced by filtering is of little or no consequence. FL1 and FL12 are still reasonable performers.

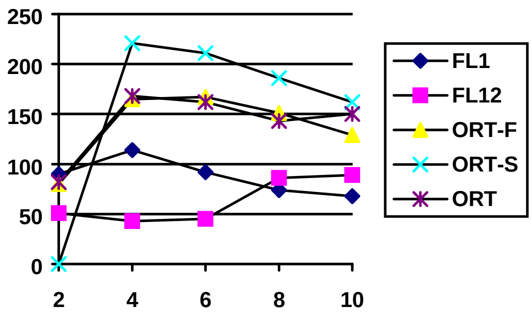

Fig 7: Mean distance of non-cluster agents - 0% noise

Figure 7 shows the increased separation obtained by using the ORT strategy with a silent mode when compared with fixed step length strategies. Although the separation achieved by ORT-S appears high, it predominantly reflects a failure to attract additional agents into the beacon area in the first place, and is a poor overall choice.

Figure 10 (at end of paper) shows some representative screen shots of the evolution of a seeded clustering sequence.

## 7 Unseeded controlled clustering

In the case of unseeded clustering, there is no beacon to use as a reference for direction of movement. The agents can therefore be simplified, because they do not require any beacon sensors. However, this means that the only source of information is the sensing of the chorusing inputs, and this might be thought to present a severe problem because the signal itself is intermittent, and the chorusing input sensor is omnidirectional. The solution to this problem is simply to use a two step klinokinesis mechanism as described in (Holland & Melhuish [1996]}.

Klinokinesis is the name given to various methods of achieving movement up or down a spatial gradient of stimulation by altering the rate and/or direction of turning as a function of the size and/or sign of some input which does not itself carry any directional information. Schöne [1980] describes an inverse klinokinesis (gradient ascent) in several types of bacteria: "Randomly distributed turnings, or jerks, interrupt the straight or slightly curved pathway of the bacterium so that it tumbles back and forth. When the bacterium enters a higher concentration of stimulant, the frequency of jerks decreases. As a consequence, the animal swims longer stretches in this direction. The sum of all the inter-jerk stretches results in a translocation up the gradient".

This technique of klinokinesis allows an agent to move up or down a gradient without being able to sense the gradient direction. The instantaneous strength of the chorusing field is all that is required, and this is already present in the agent. If the agents are in a localised cluster, then there will be a gradient of chorus signal strength which decreases with distance from the cluster; an agent which is some way from the cluster will therefore be able to move towards or away from the cluster by using klinokinesis in this field.

Early trials of unseeded clustering using the $\gamma$ mechanism and klinokinesis were only intermittently successful. Even when all agents were started at the same location, the clusters which formed were not stable even over relatively short time periods. We reasoned that analogues of the additional techniques used in seeded clustering might help; however, devising suitable implementations was not straightforward. To implement a silent mode, a measure of distance from the focus of a cluster is required, but none was directly available. In theory, it would have been possible to combine the estimate of cluster size with the sensed intensity of the chorusing signal to derive a function giving distance from the cluster, but this smacked of computational complexity and was not investigated. A simple, though partial, method is to arrange for agents to be silent when carrying out the 'moving away' behaviour; when they reach a point where they are sufficiently far from any complete cluster to stop moving away from any such cluster, they will start to chirp again. This also satisfies the requirement that agents far from the centre of any cluster should chirp, because if they do not, there can be no process of aggregation from an initially low density distribution.

The addition of a silent mode improved matters, but performance could still be very poor. A possible source was the reduced reliability of the cluster size detection mechanism in the unseeded clustering situation. Because the agents themselves are asynchronous and have stochastic elements, each agent makes frequent errors in estimating group size. An agent on its way out of a cluster could suddenly register a cluster size well below the target size, and would immediately head for the centre again, chirping and disrupting the cluster. This could be countered to some extent by tuning the filter parameter $\boldsymbol{a}$, but proved a persistent source of disturbance. An additional, but inelegant and theoretically unjustifiable 'hack' was to give each agent a type of momentum, which ensured that an agent estimating a cluster size greater than the target size would head away from the cluster for a minimum of eight steps - 4,000 time steps - before being able to change course. This was usually sufficient to remove it far enough from the cluster to avoid disrupting the cluster. However, the ultimate aim is to tune the basic mechanism and parameters so that adequate performance is achieved without this supplementary method.

Observation of the time course of unseeded clustering revealed a problem very similar to that noted by (Holland & Melhuish [1996b]) in their study of uncontrolled aggregation in these swarms of simple agents. If an initial population is dispersed over the arena, a number of small groups will form quite quickly, and each will typically be much smaller than the target size if the target size is moderately large - six or more. For a group to increase its membership, it must come within range of another group. However, since there is no factor present which will drive separate clusters together until they are within some critical distance, the only mobility process operating is a kind of stochastic drift. The larger the cluster, the smaller the rate of drift. This can result in many millions of time steps being

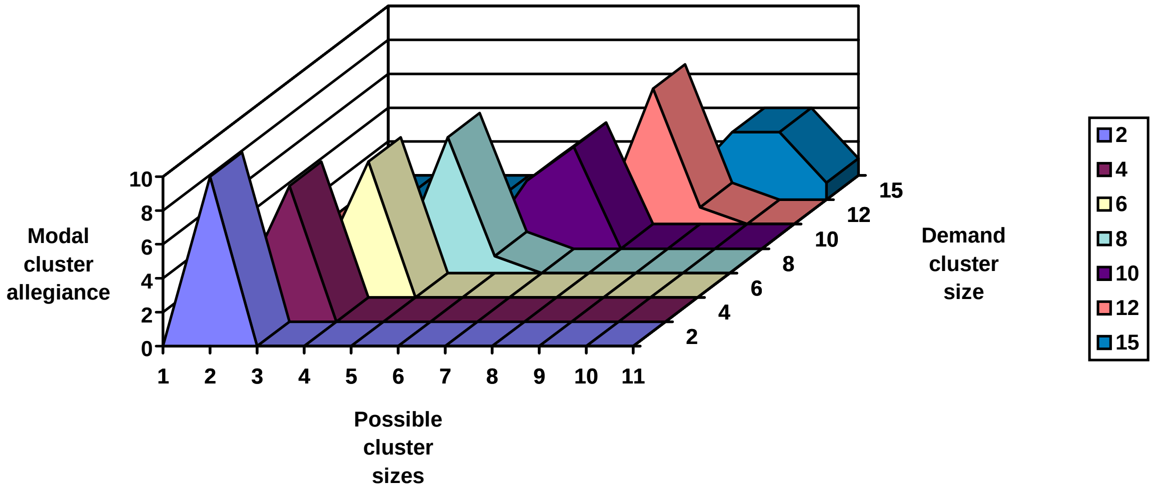

**Figure 8** **Unseeded clustering: distribution of modal cluster allegiances**

required for a cluster of the target size to be formed by the collision of smaller clusters.

A possible amelioration of this affect can be achieved by using an adaptation of the variable step length technique which is used for orthokinesis. Each agent possesses an estimate of the size of any cluster of which it is part, and a parameter which corresponds to the target cluster size. By making step length increase with increasing difference between these two factors, agents in smaller clusters become relatively more mobile than larger but still ‘incomplete’ clusters. This has two beneficial effects: the agents in smaller clusters become more dispersed than those in larger clusters, and they also move faster as a group than larger clusters. Two initially widely separated complete clusters are therefore far less likely to collide than two widely separated and very incomplete clusters; and because of the greater spatial separation of the agents in smaller incomplete clusters, large incomplete clusters will be able to attract only the nearest agent(s) of passing or approaching small clusters, enabling them to become complete and repel the remaining agents while they are still at a distance. This mechanism is called ‘error driven step variation’.

Possibly the most disappointing aspect of the performance of these unseeded clustering algorithms is their susceptibility to noise. We have not yet determined the cause of this, but we suspect that it may be due to the combination of the intermittency of the chorused signal, which is the only source of directional information for the agents, and the two-step gradient ascent algorithm which depends on the sign of the difference between two successive samples taken at discrete times. It may be possible to isolate the reason for the problem by providing the agents with better information for gradient ascent, possibly using some auxiliary data source giving an accurate and continuous gradient, and seeing if this improves noise tolerance to the levels of the seeded clustering algorithms. Fig 8 shows the effectiveness of the seedless clustering algorithm in its present state of development, using the silent mode, momentum, and error driven step variation. The noise level was set very close to zero (0.01%); the small amount of noise is technically useful for ensuring that no two sensor readings are ever equal. Agents were started at the centre of the arena; this is preferred to starting them at random positions because it reduces the early losses of agents at the absorbing boundaries, and shortens the time required for each run. When started in this way, all agents are in a single cluster which is much larger than the demand cluster size, and so they all move away from the centre in a sort of ‘big bang’; once they are sufficiently dispersed, this general movement ceases and local interactions take over. This technique ensures that a number of small groups of various sizes are formed at a reasonable distance from one anther and from the periphery. From random starting positions, it can take a great deal of computational time to form reasonably large clusters; this method results in significant time savings. Ten trials were run for each demand cluster size, and measurements of agent positions were taken after one million time steps. The agents were grouped into clusters by the simple procedure of declaring any agents within a certain minimum distance of one another to be in the same cluster.

Analysing and displaying the results presents some problems. The mean cluster size is no longer a good measure, because in a situation where the demand size is 12, a perfect outcome might be a cluster of 12 and a cluster of 3, giving a mean size of 7.5 for a perfect result. Taking the largest cluster size is also unsatisfactory; for example, with a target size of 3, a distribution of 4, 3, 3, 3, 1, 1, would be inadequately described by 4. We therefore decided to use the cluster size in which agents were most frequently found, which we called the modal cluster allegiance. The graph of

Figure 8 shows the distribution of modal cluster allegiances across the ten trials for each demand cluster size, and reveals the performance of the algorithm to be reasonably successful.

Figure 11 (at end of paper) shows some representative screen shots of the evolution of an unseeded clustering sequence.

## 8 Transferring the mechanism to a multirobot system

One axiom of behaviour based robotics has not changed in the last twelve years: simulation is never adequate, and ideas should be tested on a real robot system at the first opportunity. One of the research platforms at the IAS Lab is the Ubot, a 10-inch diameter autonomous mobile robot with differential drive, closed loop motor control, and a Motorola 68332 processor. Fifteen of these are under construction; they will be used in a variety of collective robot projects.

A chirping mechanism has been developed, and is shown in Figure 9 mounted on a Ubot. It consists of a Polaroid ultrasonic transducer mounted horizontally, facing down onto a shaped diffuser. When the transducer transmits, the diffuser reflects the highly directional pulse through a right angle to form an omnidirectional pulse parallel to the floor; when it receives, it gathers pulses from other robots and guides them to the transducer for reception. This arrangement works satisfactorily on the bench. Infra-red transmitters and receivers with an effective range of 4m have also been developed; these will be used as beacons and beacon sensors in seeded clustering experiments. These technologies, and the seeded and unseeded clustering algorithms described in this paper, will shortly be evaluated on a group of Ubots. We are not expecting immediate success even with the seeded algorithms, because the noise levels of all signals will almost certainly be at least as high (3% to 6%) as the levels leading to failure in simulation.

## 6 Conclusions

We have shown in simulation that it is possible to use the biologically inspired principle of synchronous chorusing for the control of group size in very simple multiple mobile agent systems. Several related algorithms give reasonable performance when the group is required to form at a particular point which is the source of a signal attenuating with distance (seeded clustering); performance is disrupted by noise levels of 6%. When there is no focal point (unseeded clustering) the algorithms require significant and perhaps unjustifiable modification before any useful degree of control is obtained; performance is disrupted by even small amounts of noise.

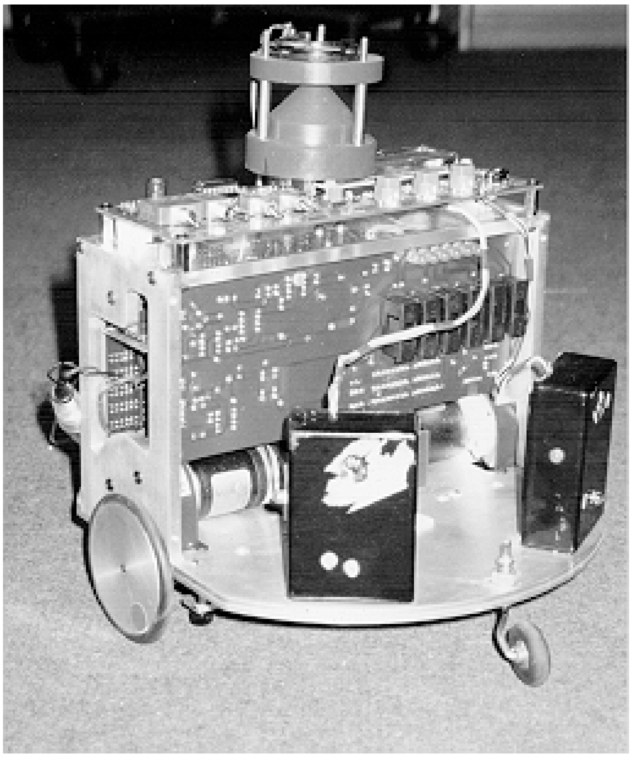

**Figure 9: a Ubot with the prototype chirp mechanism**

## Acknowledgment

This work was partially funded by Royal Mail Strategic Headquarters.

## References

Alexander R.D. [1975] Natural selection and specialised chorusing behaviour in acoustical insects. ***Insects, Science, and Society***, ed. D. Pimentel, Academic Press

Beckers R., Holland O.E., and Deneubourg J-L. [1994], From local actions to global tasks: stigmergy and collective robotics, in R. Brooks and P. Maes, eds, Artificial Life IV, Proceedings of the Fourth International Workshop on the synthesis and simulation of living systems, MIT Press

Endler J.A. [1993] Some general comments on the evolution and design of animal communication systems. ***Phil. Trans. Roy. Soc. London B*** 340, 215-225

Ewing A.W. [1989] ***Arthropod Bioacoustics: Neurobiology and Behaviour***. Edinburgh University Press

Fontan M.S. & Mataric M.J. [1996] A study of territoriality: the role of critical mass in adaptive task division. Fourth International Conference on the Simulation of Adaptive Behaviour

Greenfield M.D. [1994] Synchronous and alternating choruses in Insects and Anurans: common mechanisms and diverse functions. ***Amer. Zool.*** 34: 605-615

Holland O.E. & Melhuish C.R. [1996a] Some adaptive movements of animats with single symmetrical sensors. Fourth International Conference on the Simulation of Adaptive Behaviour

Holland O.E. & Melhuish C.R. [1996b] Getting the most from the least: lessons for the nanoscale from minimal mobile agents. Artificial Life 5, Nara, Japan

Holland O.E. & Melhuish C.R. (1997) An interactive method for controlling group size in multiple mobile robot systems. International Conference on Advanced Robotics, Monterey

Kube C.R. and Zhang H. [1992], Collective robot intelligence, in 'From animals to animats: Second International Conference on Simulation of Adaptive Behaviour', pp. 460-468

Mataric M.J. [1992], Designing emergent behaviours: from local interactions to collective intelligence, in 'From animals to animats: Second International Conference on Simulation of Adaptive Behaviour'

Ryan M.J., Tuttle M.D. & Taft L.K. [1981] The costs and benefits of frog chorusing behaviour. ***Behaviour Ecology and Sociobiology*** 8: 273-278

Schöne H. [1980] ***Spatial orientation: the spatial control of behaviour in animals and man***. Princeton University Press.

Sebeok T.A.(Ed) [1977] How animals communicate

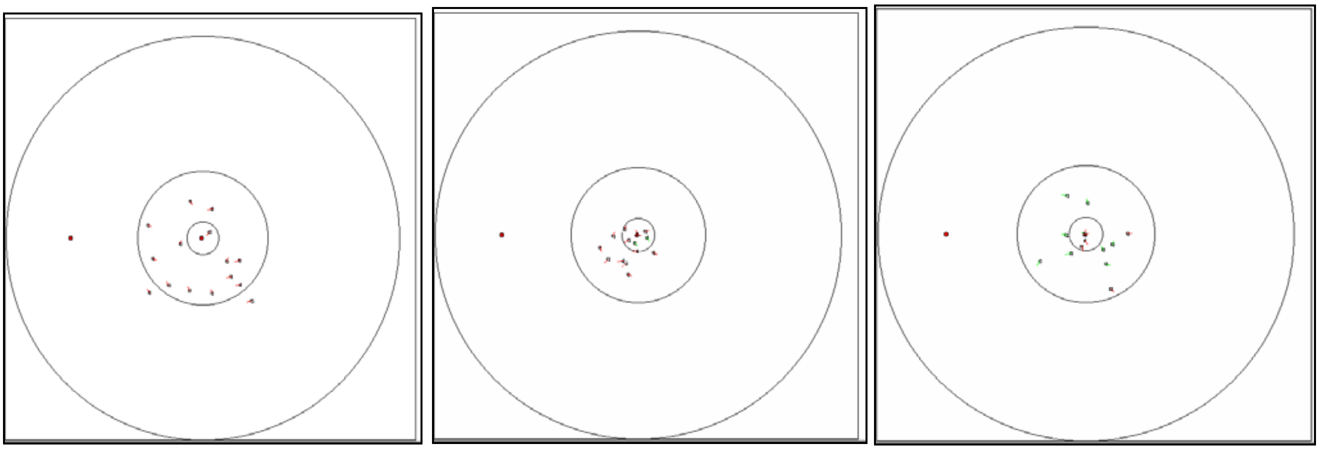

**Figure 10: Stages in the development of seeded clustering (demand cluster size 4). The inner circle defines the cluster boundary (radius 25 units). Left: shortly after the start. Centre: after 10,000 time steps there are 8 agents in the cluster. Right: after 50,000 time steps there are 5 agents in the cluster.**

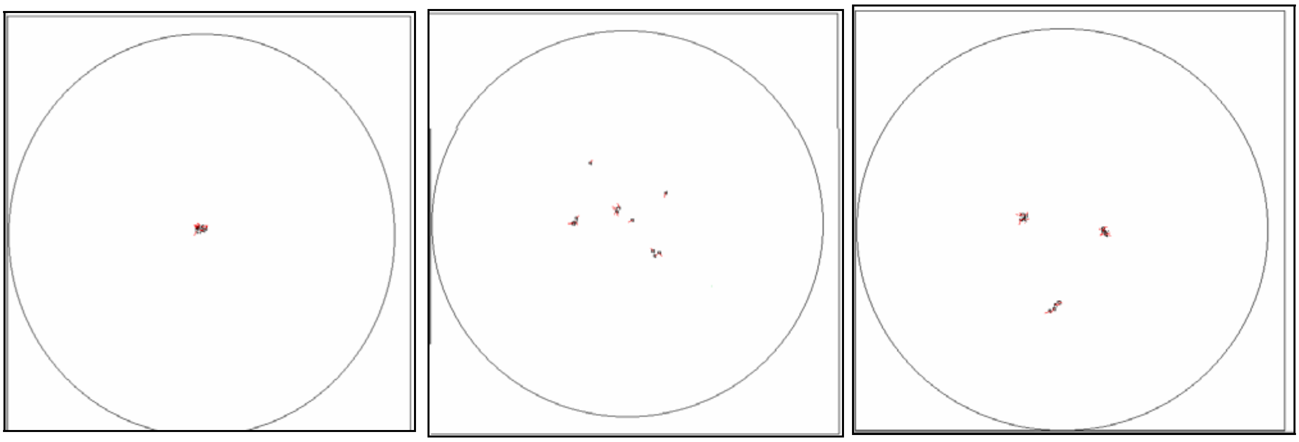

**Figure 11: Stages in the development of unseeded clustering (demand cluster size 10). Left: all agents are started at the centre. Centre: after 250,000 time steps the cluster sizes are 6, 3, 3, 1, 1, 1. Right: at the end of the run, the cluster sizes are 6, 5, and 4.**