Minimalist Collective Gradient Ascent with Real Robots: Implementing Secondary Swarming

Chris Melhuish, Jason Welsby
Intelligent Autonomous Systems (Engineering) Laboratory
Faculty of Engineering, University of the West of England
Coldharbour Lane, Frenchay, Bristol BS16 1QY

Phil Greenway, W.A.Wright
BAE SYSTEMS (ATC)
FPC 267, PO Box 5, Filton,
Bristol BS34 7QW

Contact: chris.melhuish@uwe.ac.uk

Abstract: This paper addresses and highlights some of the problems facing designers and those who engineer small scale robots in the future. It specifically looks at the problem associated with a group of small-scale robots ascending a gradient field. In particular the performance of an individual, minimalist robot can be improved when a group of similarly limited robots is employed; being a member of a collective confers benefit to the individual. This paper reports on the implementation of earlier simulation work with a group of real 'blimp' robots, with a severely restricted payload, demonstrating that spatial integrity of a group of agents around a target can be improved when employing the mechanism of secondary swarming.

1. Introduction

This paper describes the implementation of a set of collective minimalist sensing and control algorithms [Melhuish and Holland 1996] on a group of *real* autonomous robots, which are severely limited in their computation, communication and sensing capabilities. This paper reports on the utility of algorithms that have previously only been demonstrated in simulation [Melhuish, 1999]. Here it is shown that it is possible to use these ideas to control a group (or flock) of helium balloon robots which have a maximum payload of 93g and limited communications.

The ability of a group of robots to reliably perform their task *collectively* employing limited computation, sensing and communication offers many potential advantages. For instance, a collective group has built-in redundancy, in that it can withstand a certain amount of agent loss and still achieve its goal, whereas a solitary complex robot may be disabled by the loss of a single sub-system, making the task impossible. Recently, engineers have drawn on such lessons from nature and developed groups of simple autonomous mobile robots, which use simple (or minimalist) rules to act co-operatively in the pursuit of a shared goal. The roots of this minimalism spring from the study of what are sometimes called 'lower animals', where 'lower' relates to the degree of behavioural sophistication. Through what we *perceive* to be simple behaviours, collections of such animals can perform tasks that transcend the capability of the individual. For example, studies of social insects have shown that individuals, 'limited' in their ability, can collectively achieve remarkable feats and that this appears to be achieved without recourse to many of the aspects often considered necessary for "intelligent" behaviour [Melhuish, 1999]. Recently, great interest has arisen in the area of collective systems and many researchers, using these techniques have implemented multiple robotic systems for example; [Kube & Zhang (1992), Mataric (1994), Beckers, Holland & Deneubourg (1994), Kelly & Keating (1996), Melhuish, (1999)], Gulliford (1999).

With the development of miniaturisation it is reasonable to speculate how small robot machines can be controlled and co-ordinated. At the small scale there are considerable practical limitations. Small-scale robots (SSRs) will be seriously constrained in power, communications, sensing, actuation and computation furthermore, since energy sources cannot easily be carried around. These limitations are compounded by the assumption that at small scales the working environment can become extremely hostile. In these circumstances it is argued that the ability of a single robot to survive for any length of time will be doubtful and multiple robot solutions will be required. The question therefore is '*How can one get a lot of dumb robots to do something smart and in a fashion that is more robust than a single robot'?*

This paper builds on extensions to reported work on collective minimalist locomotion strategies by Holland and Melhuish (1996). This work, inspired by the locomotion of bacterium, showed that single and multiple agents, with minimal computational complexity, could move appropriately in relation to energy gradients. Collective motion was achieved in response to static and moving targets. Here strategies were developed such as *pseudo-swarming* which involved no interaction between agents and *secondary swa*rming where interaction between agents is governed by the transmission and detection of a simple secondary field generated by the agents. The work showed that performance improvements, such as in the task of homing on a beacon, could be achieved by providing the simulated robots with additional capabilities. These were inspired by observations of social insects that employ a form of recruitment by generating an acoustic signal when they have found a quality food source. For example, once a leaf cutter ant has found an appropriate leaf it is able to signal its find to other ants, which lie outside the proximal detection range. The solution employs a secondary homing source. That is, each agent is allowed to be the source (and receiver) of minimalist sensory information. It was argued that an agent could 'home-in' on the target if it could detect it. On detection of the target it would also generate a secondary 'localised' field. This secondary field could be of considerably less 'broadcasting power' than the primary target (in fact the power requirements of this secondary field can be reduced further by the use of synchronous short broadcast bursts [Melhuish et al.1998a; Melhuish 2001], An agent unable to detect the primary target (the signal being less than some threshold) would attempt to use the same locomotion strategy but employ the summed secondary fields, generated by those agents which can detect the primary target, rather than the primary target itself.

Employing such strategy provided a cohesive effect for a moving swarm Melhuish [1999], where the agents nearer the target occluded their neighbours further away, even though those ‘near’ agents were not a fixed set of agents since they consisted of an ever-changing pool of agents within the swarm. It was also shown that, for the model employed, as more robots found the target and became secondary field sources, the range of the combined secondary field increased, resulting in the recruitment of more robots.

Current implementations of collective movement with physical robots is often constrained to two dimensions [e.g. Kelly & Keating, 1996]. In contrast to this there have been many demonstrations of this behaviour in three-dimensional simulations, possibly the best known being that of Reynolds (1987). Here, a group of real robots that are not constrained to two-dimensions (although the robots are not given full vertical freedom within the constrains of the laboratory) are described. The robots utilise helium balloons (blimps), and are therefore constrained to a limited payload in which to fit the propulsion, communication and localisation systems (see section 3). The design of a suitable computation, sensing and power system are described in section 3. The details of the control algorithms are also described here.

A description of how secondary swarming was implemented on these platforms is then described in section 4. Here it is demonstrated that there is an improvement in the performance; any single robot in the group performs better than a robot acting completely on its own. These results also show that, to a degree, the problems caused by hardware variability can be decreased significantly if the robots work together. This is based on a comparison of the performance of the system when implementing two separate strategies, pseudoswarming and secondary swarming, to enable the robots to home on to a static beacon. Pseudoswarming is where the robots individually home in on the beacon with no interaction between each other but still give the impression of collective behaviour.

2. Minimalist Strategies

The payload capacity of the blimps imposes severe constraints on the level of sensing, communication computation and power available to the blimps. A minimalist approach to sensing and actuation provides a number of advantages. This approach is now discussed further. From a sensing perspective it can be argued that one simple sensor is more minimalist than two simple sensors [Melhuish, 1999]. However, to obtain useful information from a single sensor requires two readings to be taken at different times (i.e. a two-instant mechanism [Feinleib 1980] and this requires some form of simple memory. Further, the time difference between sensor readings leads to periods of inactivity in the propulsion system. Alternatively, a system employing two sensors that takes simultaneous readings from both sensors (i.e. a one-instant mechanism) can continually adjusts the trajectory according to the difference and requires no memory. Consequently, this one-instant mechanism allows smooth movement of the robot and does not require the use of memory. Furthermore, that addition of a second sensor does not unduly add to the complexity of the system.

From a computation point of view Holland (1996) lists two broad classes of motion, impulsive and smooth. An impulsive system releases energy in bursts using intermittent, or bang-bang, control, whereas a smooth system releases energy continuously using constant control. The large inertial component of the blimp requires constant control to maintain stability and this favours a smooth movement strategy. Furthermore, such a strategy easily integrates with the one-instant sensing. Computation was kept to a minimal level using the sensed signal to implement a binary choice between the left and right motors. The canonical form of the locomotion algorithm used to move a blimp toward a simple beacon used a gradient ascent algorithm and is shown below in Figure 1. The algorithm can be considered as the conflation of sensing, processing and action and was used as the basis for all movement algorithms implemented

```
If signal is above noise threshold and below max threshold:
        If Left signal > Right signal
            Slow the left motor (turn left)
        Else
            Slow the right motor (turn right)

        // and dealing with extreme conditions:
        If signal is above max threshold:
            Back away (both motors reverse)
        Else If signal is below noise threshold:
            Stop (randomly move with air currents)
```

Figure 1. Beacon Homing Behaviour (Canonical Form). Note 'signal' here refers to the received beacon signal

In order for the blimps to act collectively they must be given the ability to affect each other under certain conditions. To this end the collective algorithms used indirect communication in the form of a transmitted secondary field to enable group interaction. Here, to all intents and purposes, the robots see this secondary field as an extension of the primary beacon field. Secondary swarming was employed by the judicious switching of the signal source. Each robot was allowed to generate its own field (as described above). Essentially, if a robot can detect the beacon it will employ the minimalist algorithm above. However, if it fails to detect the beacon it will attempt to climb the gradient field generated by a robot that can detect the beacon (see Figure 2). Full implementation details are described later in section 4.

```
If robot can detect the beacon:
    Ascend gradient field generated by the beacon
Else
Ascend gradient field generated by another robot/s
```

Figure 2. Exploiting the secondary field

3. Materials and Methods

3.1 Construction of the Robot Blimps

The physical implementation of the robot uses a lighter than air vehicle (LTAV) in the form of a helium filled balloon (blimp). The blimp consists of a two-panel metalised nylon envelope 96.5cm in diameter uninflated, which once filled has the shape of a squashed sphere (figure 3) 0.75m in diameter. The volume of the envelope is approximately $0.3m^3$ giving a net lifting capacity of 93g when filled with Helium/Air balloon gas, this gas is only 93% Helium making it much cheaper than pure Helium with a negligible difference in its lifting capacity. The blimps had an onboard computer, propulsion system and an infrared localisation and communication system. A lithium-ion battery that gives an operational time of 2½ hours provides the power. These are now described in the following sections.

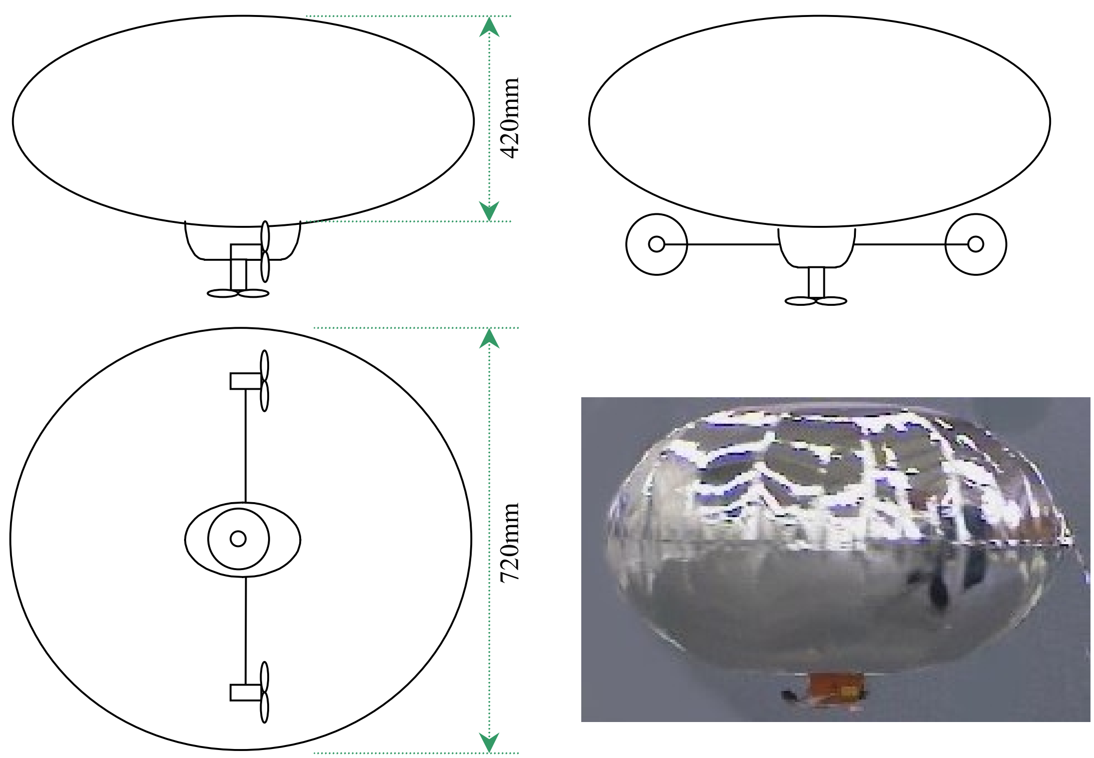

Figure 3. Physical Dimensions of the Inflated Envelope

3.1.1 Gondola

The blimp gondola was made of lightweight plastic 'blister packs' which are commonly found holding products to cardboard backing. It used the packaging from 'Mini Jaffa Cake' biscuits and provides enough space to hold all the electronics and batteries. Plastic drinking straw arms fixed to the plastic housing attached the motors to the gondola. The plastic straws were flattened at the ends, bent around the motors and heat welded to give a firm fixing point that required no glue or clips (figures 4 and 5) thus minimising weight.

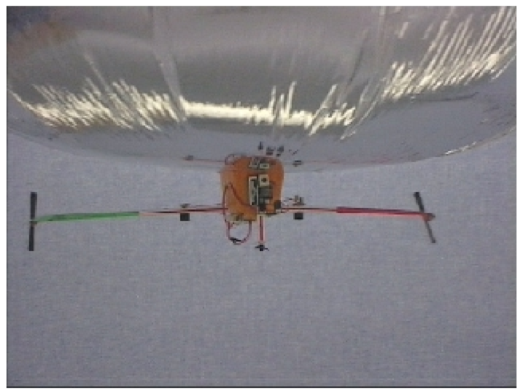

Figure 4. Gondola attached to the envelope.

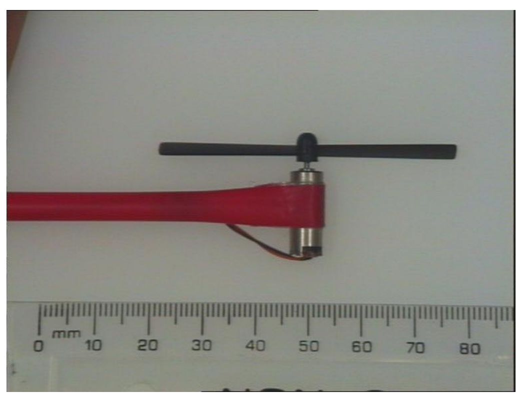

Figure 5. Fan unit mounted on plastic drinking straw arm.

3.1.2 Propulsion System

Thrusting on the blimp was achieved with three small fan units capable of supplying approximately 4g of thrust at full power. Each of these units consisted of a small DC electric motor fitted with a small plastic propeller (figure 6). The motors were the same as those used for vibrating mobile phone batteries and pagers, and are very lightweight (2g). The plastic propellers were 5cm in diameter and weigh 0.3g each. Each of the fans was controlled with a full bridge driver (h-bridge) to enable operation in both directions. Pulse-width modulation (PWM) was used to give digital speed control.

One fan was mounted at each end of an arm to provide forward, reverse and axial rotation of the blimp (figure 4). The third fan was fitted beneath the gondola to provide movement in the vertical axis. A problem with this vertical axis fan was that it induced small torque, causing the blimp to rotate in the opposite direction to the propeller. However, if only used for a short period this torque was negligible compared to movement caused by air currents, and could be overcome by using a gimballed design instead of the third fan unit. The arm holding the two thrust fans could then be tilted 90^0 either way of the horizontal to provide twice the thrust in the vertical axis. This design would have required a servomotor to move the arm, which weighed much more than the single fan unit it replaced. Thus for the simplicity the gimballed arm was not used.

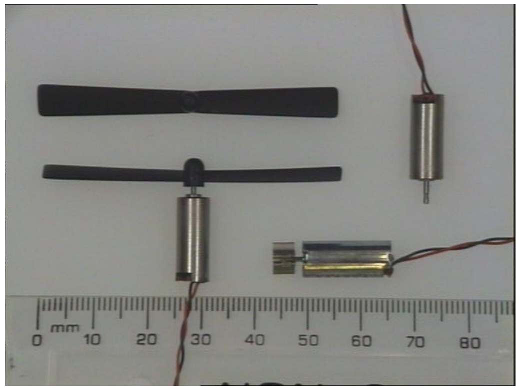

Figure 6. Motors and propellers with steel rule for scale.

3.2 Control System

The main drawback of using small blimps is the severe weight limitations placed on the design of the electronic control system: the net lift of the inflated envelope was only 93g and more than half of this (50g) was used for batteries. To save weight all the circuits were manufactured from 0.5mm fibreglass board, which is one-third the thickness of a standard circuit board. This thin board coupled with the very compact layout, and the use of surface-mount technology, reduced the weight by a considerable amount. The only problem with this compact design was that assembling it by hand was a very intricate task. Figure 7 below shows a biscuit containing all the unpopulated circuit boards.

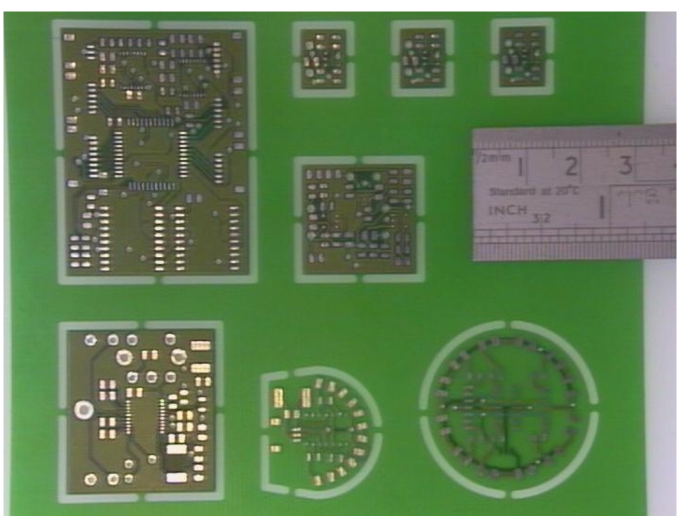

Figure 7. Biscuit of Double Sided Circuit Boards

For this design the electronics were split into distinct parts containing the digital and analogue part of the system separately in order to minimise noise coupling.

3.2.1 Digital Board

The digital circuit consisted of the power generation, motor drives, frequency generators and a microcontroller. These all operated at high frequencies producing a significant level of noise and so interference.

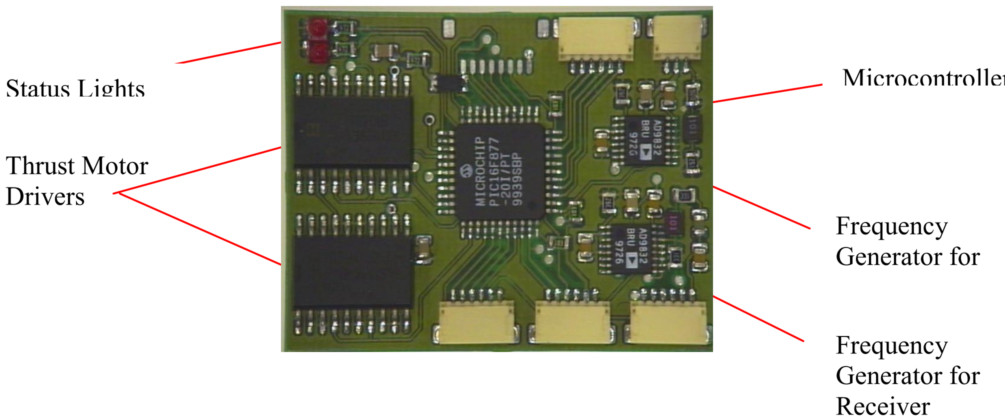

Figure 8. Control Board Top

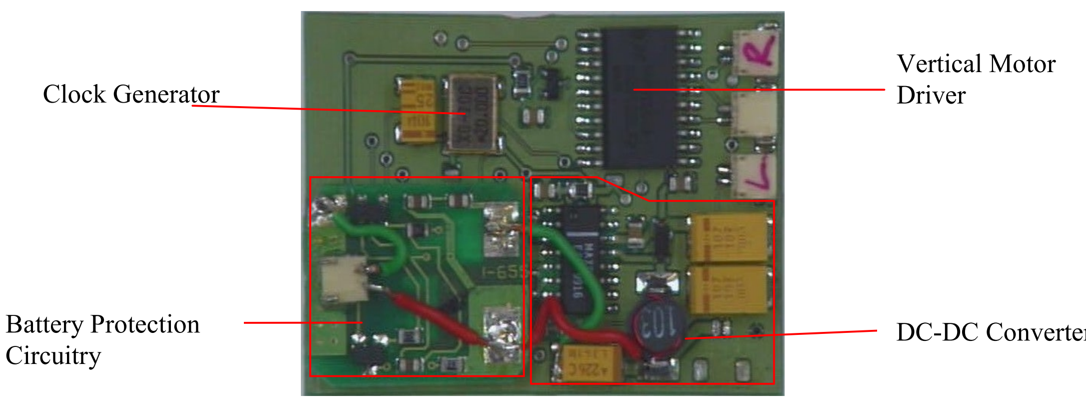

Figure 9. Control Board Bottom

Given the step-up DC-DC switching converter (figure 9) and a rechargeable battery supply 5 volts needed by the electronics the maximum current consumed by the electronics could exceed 1Amp and consequently therefore a high current switching converter was needed. The MAX489 was chosen since it has a load capability greater than 1Amp and is designed to be very low noise (4.7nV/Hz). This converter works with inputs as low as 0.8 Volts, maximising the energy that can be drawn from the batteries. However, the device operates at 300KHz and this had implications for design of the IR localisation system (as explained in section 3.3.2).

The motor drivers (figures 8 and 9) were HIP4020 full bridge drivers capable of delivering 500mA in an SO20 surface-mount package. These chips also have over-temperature and over-current cutouts to prevent overheating and damage. The use of this chip helped keep the component count to a minimum, saving weight and board space.

The frequency generators (figure 8) used by the transmitter and receiver use direct digital synthesis (DDS) technology. Theses produce near perfect sine waves with minimal external components. The ICs used were AD9832s controlled via a three wire serial interface. These came in very small (TSSOP) surface mount packages again saving weight and space. The original localisation system (see section 3.3.2) used a separate crystal for each frequency, taking up valuable space and adding unnecessary weight. However, by replacing all the crystals with a single IC, weight, board space and cost were minimised. The new design also made it easier to change and add new frequencies, via a software interface, when desired.

The heart of the control system is an 8-bit Arizona Microchip PIC 16F877 microcontroller (figure 8). This runs at 20MHz giving ample computational power. This microcontroller has 8Kbytes of onboard flash memory, which can be programmed in-circuit via a three-wire serial interface. The unit also has two PWM and three timer modules. The serial port on the PIC gave full RS232 functionality when used in conjunction with a special lead that had a built in level converter chip to provide the correct voltage matching.

3.3 Sensors

3.3.1 Ultrasonic Height Control

A lightweight ultrasonic ranging system that gave the blimps the ability to control their height to within a few centimetres with a maximum range of 3-4 m was designed. This sensing system is highly accurate (within a few millimetres) however; the dynamics of the blimp and its propulsion system did not allow accurate height control. The ranging system used a transmitter to bounces a 40KHz ultrasonic pulse off the floor, which was then picked up by the receiver to give a time of flight and hence the distance. For simplicity, and to minimise processing power, the first echo received was used to calculate the distance (i.e. the system does not check to see if the first signal was valid by averaging over a number of pulses). This system worked well over a smooth flat surface but gave reduced accuracy over surfaces such as carpet.

The accuracy and range were not only affected by the reflectivity of the ground surface but also the alignment of the receiver and transmitter. To gain maximum range and accuracy the transmitter was placed at least 50mm from and parallel to the receiver, this prevented any direct mechanical coupling between the transmitter and receiver. The system used small separate ultrasonic transducers for the receiver and transmitter, as shown in figure 10 below. This allowed short-range measurements to be made quickly since a single transducer would require a blanking period to enable it to settle after a transmission. To stop false triggers through conducted noise a time decaying threshold system was employed. Initially a high threshold was used which then decayed, to a pre-set minimum, to allow the weaker more distant echoes to be registered. To save weight the circuit boards were made from standard 3mm board, which were reduced in thickness to less than 1mm. Figure 9 shows the transmitter and receiver boards with the transducers not fitted. Note this system could also be used for forward facing collision avoidance since it can detect objects more reliably and accurately than the infrared system.

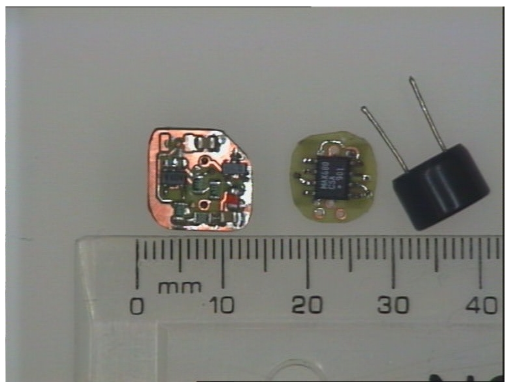

Figure 10. Ultrasonic transmitter and receiver boards with a sensor

3.3.2 Infra-Red Localisation System

Each robot is fitted with a forward facing 270^0 infrared receiver whose reception angles are indicated by the light shaded areas in the top and side views of figure 11. As can be seen in the side view, the blimp envelope obscures the receiver from above creating a large blind area above the blimp. Each blimp is also fitted with an 180^0 rear facing infrared transmitter whose dispersion angles are shown by the darker shaded areas in figure 11.

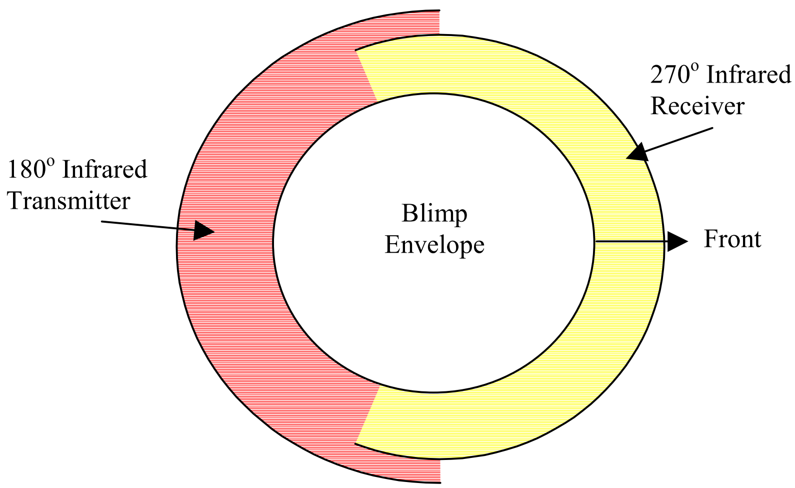

Top View

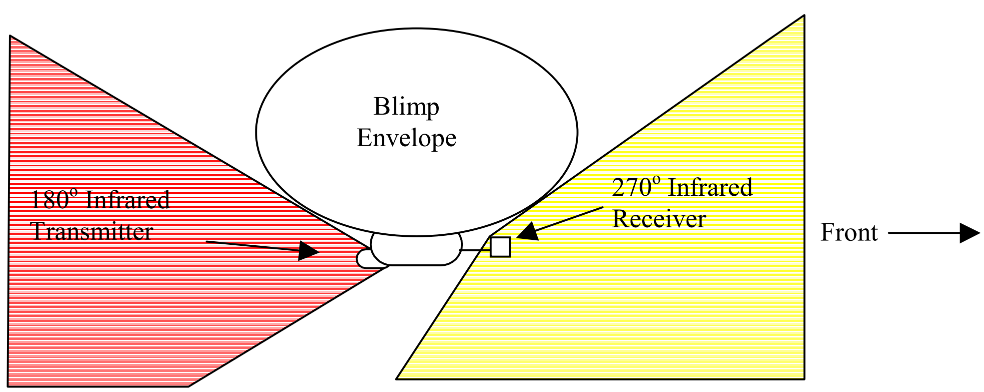

Side View

Figure 11. Showing transmission and sensing angles of the blimp

The Infrared localisation system is based on an original design by Kelly [1996]. This gave the relative position and distance of the other blimps, and allowed for low bandwidth inter-blimp communications. Three types of analogue board made up the infrared localisation system: a receiver, transmitter and two sensor boards. These boards were powered by a filtered voltage supply to minimise the noise induced by the high speed switching of the digital devices.

The heart of this system was the receiver (figure 12) which was based on a low power Phillips SA607 FM demodulator IC which had a logarithmic received signal strength indicator (RSSI). This RSSI signal was fed into an analogue to digital converter (ADC) on the control board and was then used to calculate the distance to the transmitter being scanned. The graph in figure 13 below shows the RSSI plotted against distance, and demonstrates that for the most part this relationship is 'gracefully attenuating' and so, can be easily converted to a relatively accurate distance measurement.

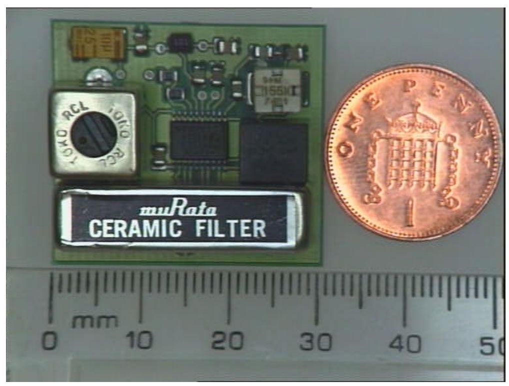

Figure 12. Receiver Board

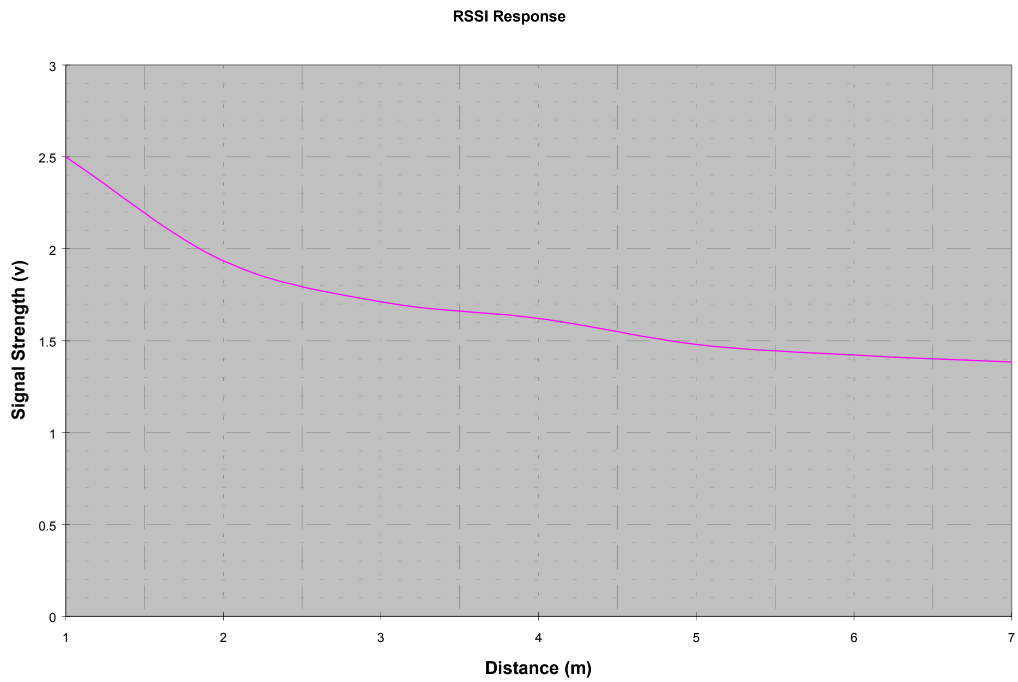

Figure 13. RSSI response (for one trial)

To give the blimps a unique signature the system uses Frequency Division Multiplexing (FDM). Here each blimp transmitted on a separate frequency selected to minimise interference between the blimps or environmental noise. At the frequencies used (200KHz to 600KHz) there were found to be a range of noise sources that interfered with the infrared system. A main source of noise, within the experimental environment, was the fluorescent lighting in the lab. These produce interference at 50Hz and higher harmonics. Fortunately there was little interference above 200KHz. Also as mentioned in section 3.2.1 the onboard switching power converter operated at ~330KHz producing a noise spike at this frequency. Figure 14 shows the background noise spectrum of the environment in which the blimps were placed. Consequently the system was tuned to operate above 200KHz in order to avoid this interference.

Figure 14. Frequency Scan from Blimp

The infrared transmitter (figure 15) consisted of a semicircle of 6 LEDs. This gives 180^0 coverage to the rear of the blimp. The LEDs used were inexpensive high power infrared emitters 'HIRL5015' which have a half power angle of 60^0. Each LED was 30^o apart in the same plane so that the light from adjacent LEDs overlaps (see figure 16), giving a radiant intensity which stays approximately constant with respect to angle minimising the lobing effect. Two LEDs were placed in series and fed with a 5V sine wave at 100mA produced using high output current Op-Amps fed from one of the frequency generators. The use of a sine wave to drive the LEDs instead of a square wave, as in Kelly's design, gave a clean output with few harmonics. By using these LEDs in conjunction with the sine wave it was possible to achieve a range in excess of 20m, nearly twice that of Kelly's original design.

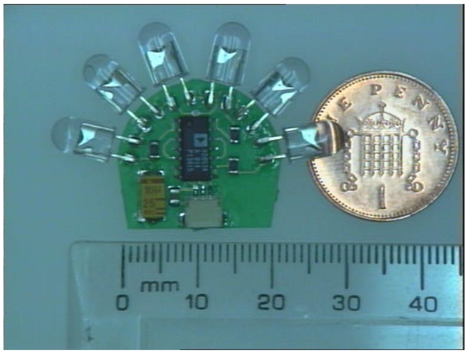

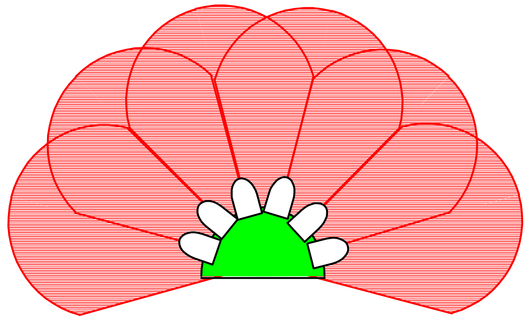

Figure 15. 180 degree IR Transmitter Module **Figure 16.** Coverage Pattern from Transmitter Module

As discussed, each blimp transmitted its own unique signal so that it was possible to differentiate between the blimps. Each blimp scanned all other frequencies in use except its own. Scanning its own signal would result in a very high RSSI signal due to the close proximity of the transmitter to the receiver.

Each blimp was fitted with three forward-facing infrared sensors (figure 17) arranged as a pyramid with its point facing forward to give an approximate hemispherical field of view. Each of the infrared sensors can be selected through an 8:1 analogue multiplexer, with the possibility of adding a further five sensors to expand the system if required. The relative position of any blimp being scanned can be calculated by comparing the data received from each sensor.

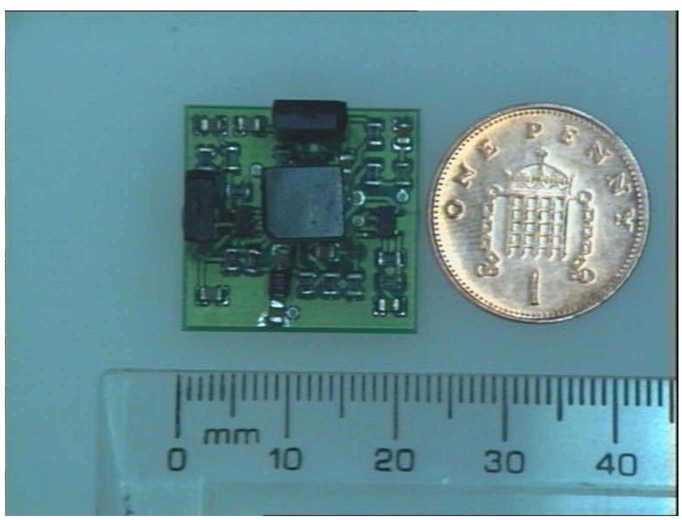

Figure 17. Three Way IR Sensor Board

3.3.3 Batteries

Battery choice was mainly influenced by the gravimetric energy density available): the lighter the battery for a given capacity the better (see the schematic shown in figure 18). It was decided to use a rechargeable (secondary) technology, because the amount of experimental run time would incur massive costs if a non-rechargeable (primary) battery technology were used. Three types of rechargeable battery were considered: Nickel Cadmium (Ni-Cad), Nickel Metal Hydride (Ni-MiH) and Lithium Ion (Li-Ion).

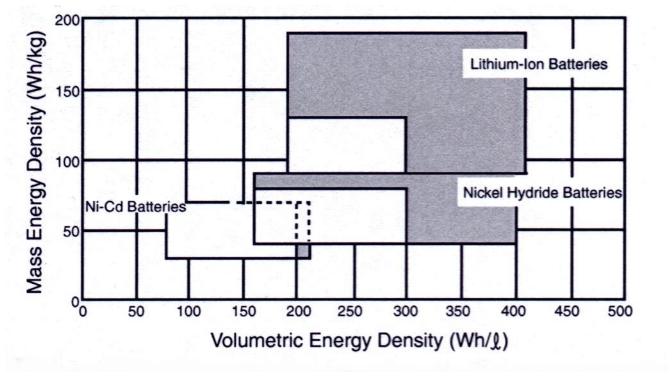

Figure 18. Battery Technology Comparison, Panasonic Corporation

As illustrated Li-Ion batteries have superior gravimetric and volumetric energy densities. After much searching it was found that the most appropriate and readily available battery was Sony's Minidisc Li-Ion battery (LIP-12B). This gives 1500mAh at 3.6V for a weight of 50g. The internal cell itself (figure 19) only weighs 39g, the remainder being made up by the protection circuitry and outer casing. The weight of the battery pack was minimised by discarding the outer plastic casing which gave a weight saving of 8g. The protection circuit was incorporated onto the control board (figure 9) with a small lightweight socket to plug in the battery cell. A further protection circuit was mounted on a charging adapter (figure 19) to allow the battery to be recharged with the safety mechanism in place. The use of this new battery gave a huge increase in the operational time of the blimp (from 10 minutes using two Duracell Ultra batteries to 2.5 hours). The only drawback with the new batteries was that the payload was pushed close to the limit, allowing only a further 8g of to be added.

Figure 19. Battery Cell and Charging Adapter

3.4 Software

The algorithms describing the behaviour of the blimps are written in the 'C' programming language. Once written, the program code was compiled using the 'Hi-Tech' C compiler for the PIC microprocessor and then downloaded to the control board via the 'MP-LAB' development environment. The use of 'C' made program development much faster and easier than writing in PIC assembly language. However, care had to be taken to ensure that the compiled code was small enough to fit into the 8Kbytes of memory available. A collection of low-level control routines provided a set of useful functions for the implementation of the behaviours.

The control board was equipped with a serial interface which allowed data to be uploaded to the PIC or downloaded to the host PC, this together with the fact that the PIC can write to its program memory while running allowed a bootloader to be implemented. The bootloader was the first part of the software system to run when the PIC is started. It checked to see if program had already been downloaded and goes on to execute the program otherwise it waited for a new program to be downloaded. Once the PIC has been programmed with a bootloader a new program can be uploaded via the RS232 port on any PC. This bootloader method negates the need for the MP-LAB development environment except when programming the PIC with the bootloader for the first time. Caution was required when compiling program code to ensure that it did not overwrite the program memory occupied by the bootloader. The RS232 interface also allowed information to be gathered by the PC, which was invaluable when debugging software or testing code for the first time.

4. Experiment Details

Using the robots a set of experiments were undertaken to discover if a collective strategy can increase the performance of a group of robots when homing on a static beacon. Two different strategies were implemented and compared; pseudoswarming (when robots act as completely independent agents) and secondary swarming.

When pseudoswarming, the robots individually home in on the beacon with no interaction between each other. However, they can still give the impression of collective behaviour as they move toward the beacon. With a small change to the algorithm the robots can be made to interact with each other changing the way in which they home on the beacon.

The experiments were conducted in the experimental arena with the beacon fixed in one corner at a set height of 1.5m and the robots started in the opposite corner (see figure 20 below) to give a separation distance of approximately 10m, (just within the blimps maximum sensing range.) Four robots were started together and their positions recorded at one second intervals by the overhead camera system. A time limit of two minutes was placed upon each trial to limit the amount of data collected, it was expected that on a good run a robot would reach the beacon within 30 seconds if it could initially detect the beacon and did not lose its signal.

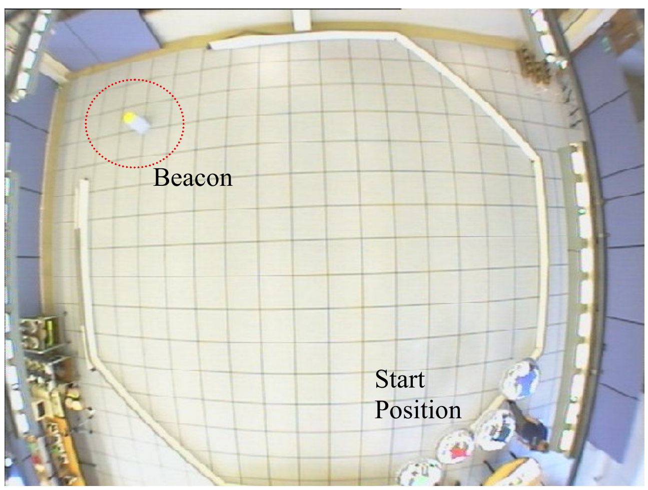

Figure 20. Showing the beacon and starting positions of the blimps

4.1 Performance Metrics

To assess the performance of the algorithms three metrics were employed: time taken for all blimps to reach the beacon, average distance of all the blimps from the beacon as a function of time and success rate of reaching the beacon expressed as a percentage.

4.2 Experiment 1: Individual Homing on a Static Beacon

The Pseudoswarming experiments aim to show how well a group of four robots home in on a beacon using purely individualistic strategies.

4.2.1 Experimental Details

The pseudoswarming algorithm

The following algorithm was implemented on each robot:

If signal is above noise threshold and below max threshold:

If Left signal > Right signal

Slow the left motor (turn left)

Else

Slow the right motor (turn right)

If signal is above max threshold:

Back away (both motors reverse)

Else If signal is below noise threshold:

Stop (randomly move with air currents)

4.2.2 Results

For this experiment eight separate trials were conducted the results of which are shown in figure 21. In order to show the homing process in action the first trial has been examined in more detail with the aid of still images from the overhead camera. In all 16 frames (figure 21 A-P) are shown covering the first 90 seconds of the experiment.

The blimps were started in one corner of the arena as shown in frame A. When the beacon was switched on the blue blimp acquired the signal immediately and started homing in. This blue blimp moved away from the others quite quickly as illustrated in frames C – F and manages to reach the beacon in 25 seconds (frame H). The other blimps spent the first 20 seconds (frames A–E) ‘milling’ around by the start unable to detect a signal from the beacon. 22 seconds into the experiment (frame F) the green blimp wanders within sensing range of the beacon and begins homing in leaving the remaining two blimps lost. The green blimp takes a further 10 seconds (frames H-K) to reach the beacon during which time the red and black blimps remain out of range. The red and black blimps continue to wander around the area they started in until at time 62 (frame M) the black blimp drifts away toward the centre of the arena and closer to the beacon. The black blimp continues its drift toward the centre for another 14 seconds at which point (frame O) it senses the beacon and begins homing reaching its target 12 seconds later (frame P).

The red blimp continues to wander just outside the beacon range for the remainder of the experiment.

Frame A 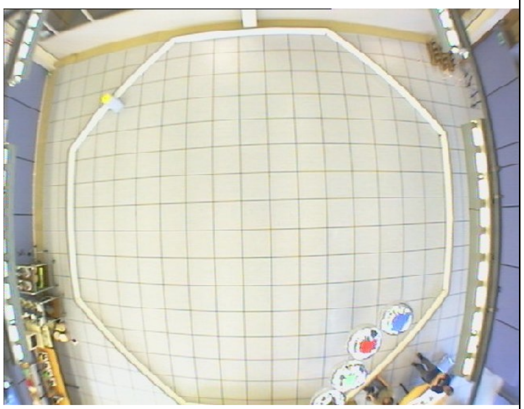	Frame B 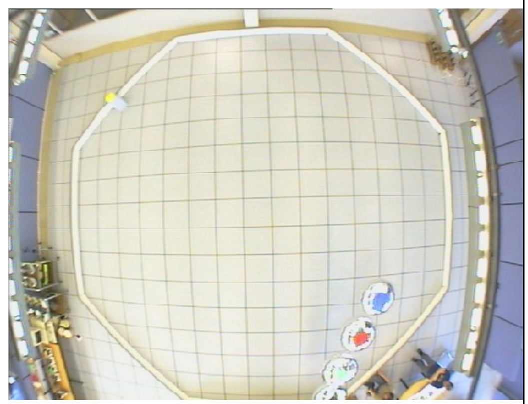
Frame C 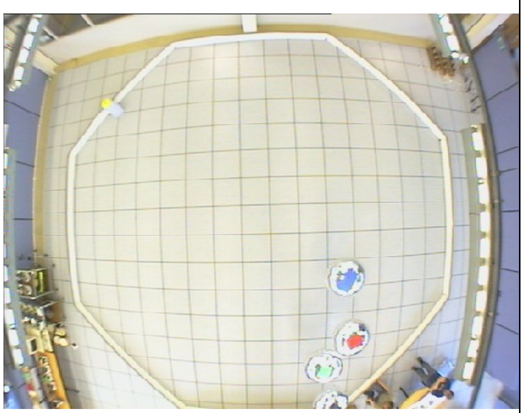	Frame D 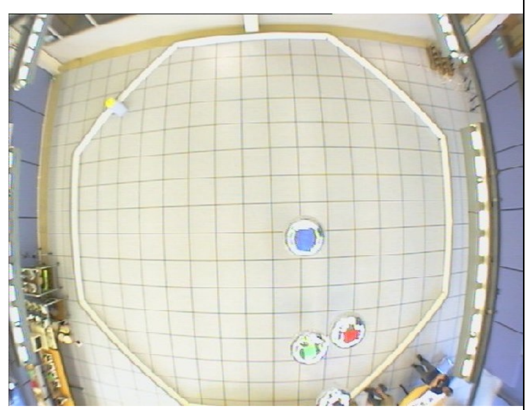
Frame E 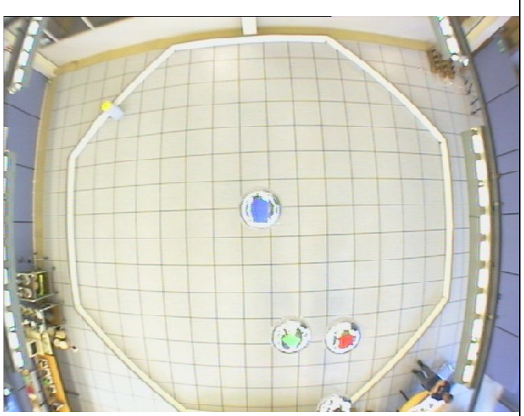	Frame F 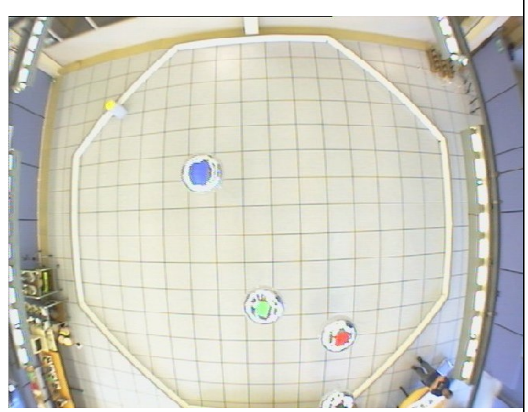

Frame G 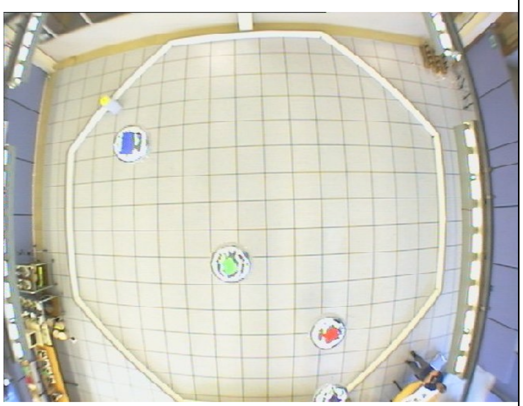	Frame H 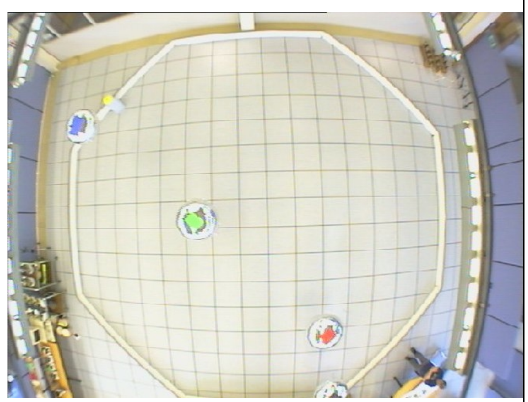
Frame I 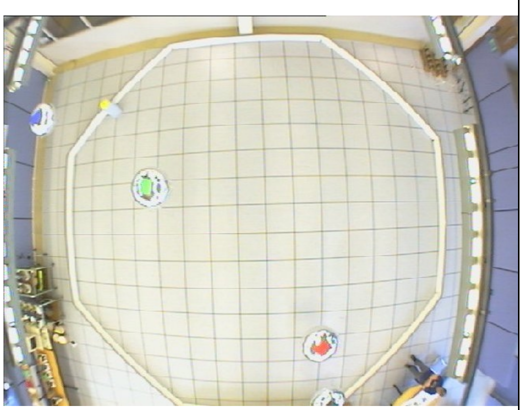	Frame J 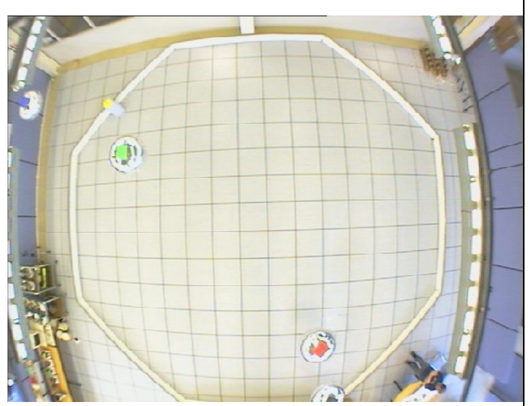

Frame K

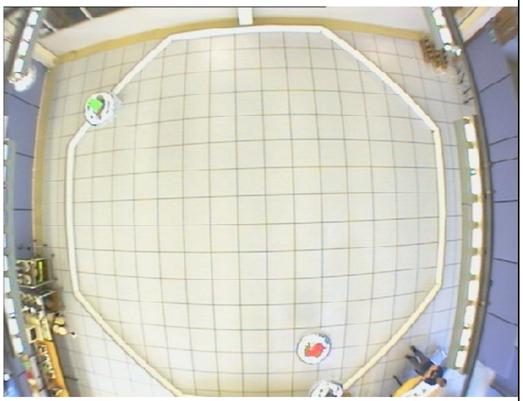

Frame L

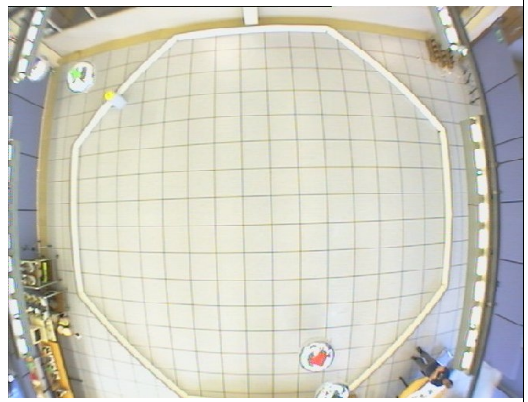

Frame M

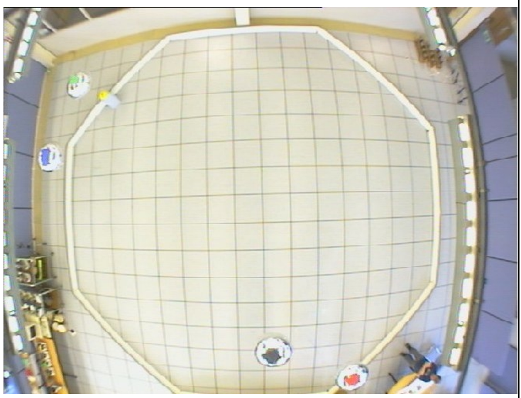

Frame N

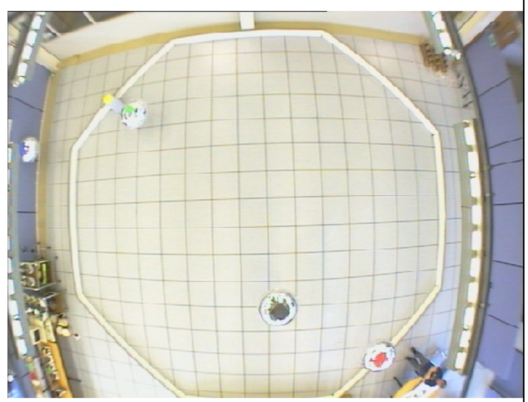

Frame O

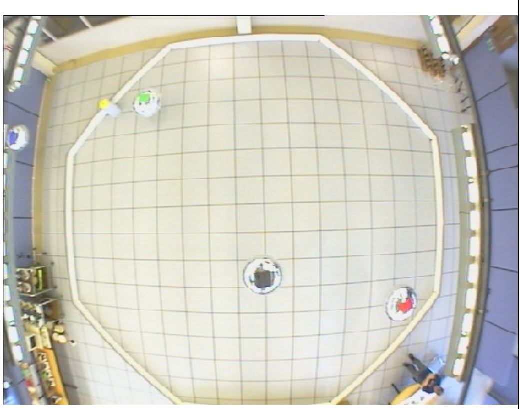

Frame P

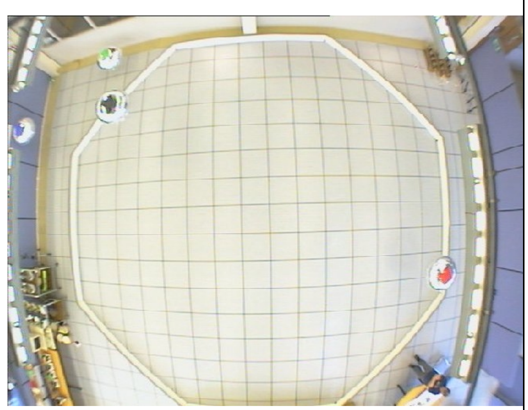

Figure 21. Sequence of frames showing individual homing on a static beacon

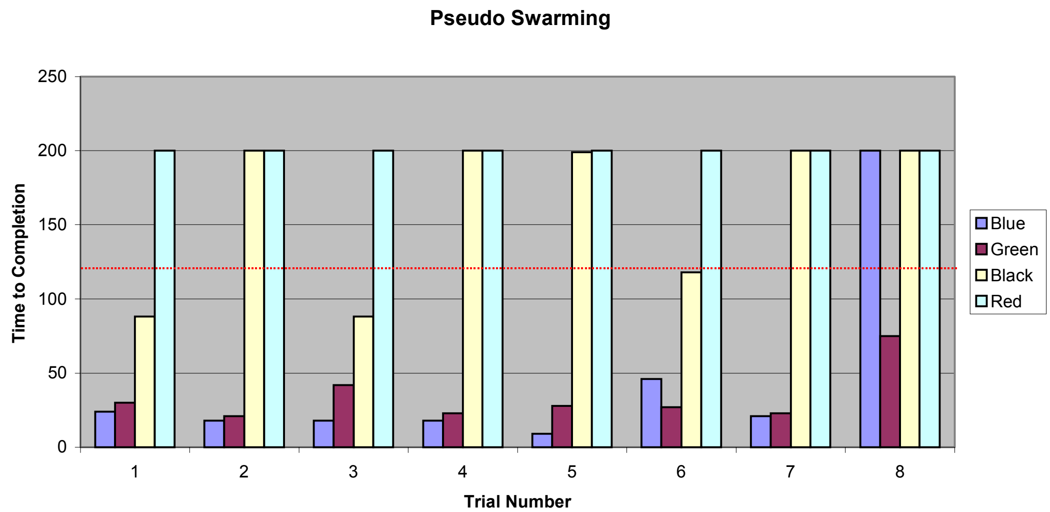

Figure 22. Results for Pseudo Swarming

The graph in figure 22 gives the results for all eight trials conducted and shows the time taken by each robot to reach the beacon. The horizontal red dotted line indicates the two-minute time limit at time 120s. It can be seen quite clearly that the blue and green blimps perform far better than the other two, on average making it to the beacon in 30 seconds or less. However, in the last trial the blue blimp fails to reach the beacon altogether and wanders off in the opposite direction. The red blimp consistently fails to reach its target destination in every trial and the black blimp only fares marginally better reaching its target 3 times out of the 8 trials. From the results shown in figure 22 it can be seen that none of the trials actually complete within the given time limit of two minutes. However, in all of the trials at least one robot makes it to the beacon.

4.2.3 Conclusions

The results show that this 'individualistic' algorithm does not produce a reliable homing behaviour for a small group of robots. After close examination it was found that the poor performance of some of the robots, and hence the experiments, was due to differences in the sensor ranges. The blue and green blimps had greater sensing distances than those of the other two robots.

Another problem, that hindered the homing process, is that of occlusion where one robot blocks the beacon signal to another robot. This beacon occlusion played a significant part at the beginning of each experiment as the blimps were close to one another and could easily obscure the signal, the effect was seen to decrease with increased spatial separation of the blimps. The problem would arise again when a blimp neared the beacon, the closer it was the larger the obstruction it posed.

4.3 Experiment 2: Collectively Homing on a Static Beacon

Secondary attempts to improve on the individualistic pseudoswarming strategy by the addition of robot transmissions when a robot can see the beacon. This effectively increases the range of the beacon and also counteracts the occlusion problem.

4.3.1 Experimental Details

The secondary swarming algorithm was tested under the same conditions as the previous pseudo swarming experiments. The same robots were used and the only changes that made were to the algorithm are detailed below:

If robot can detect the beacon:

Ascend gradient field generated by the beacon and transmit a secondary field

Else

Ascend strongest gradient field generated by another robot

If *signal* is above noise threshold and below max threshold:

If Left signal > Right signal

Slow the right motor (turn left)

Else

Slow the left motor (turn right)

If *signal* is above max threshold:

Back away (both motors reverse)

Else If *signal* is below noise threshold:

Stop (randomly move with air currents)

As in the previous experiment eight trials were run with a time limit of two minutes, the results of these trials are shown in figure 23.

4.3.2 Results

From the results shown in figure 23 it can be seen that for each trial all robots reach their objective within the given time limit. This indicates therefore, that there is a significant advantage in using collective strategies over purely individualistic techniques when homing on a static beacon.

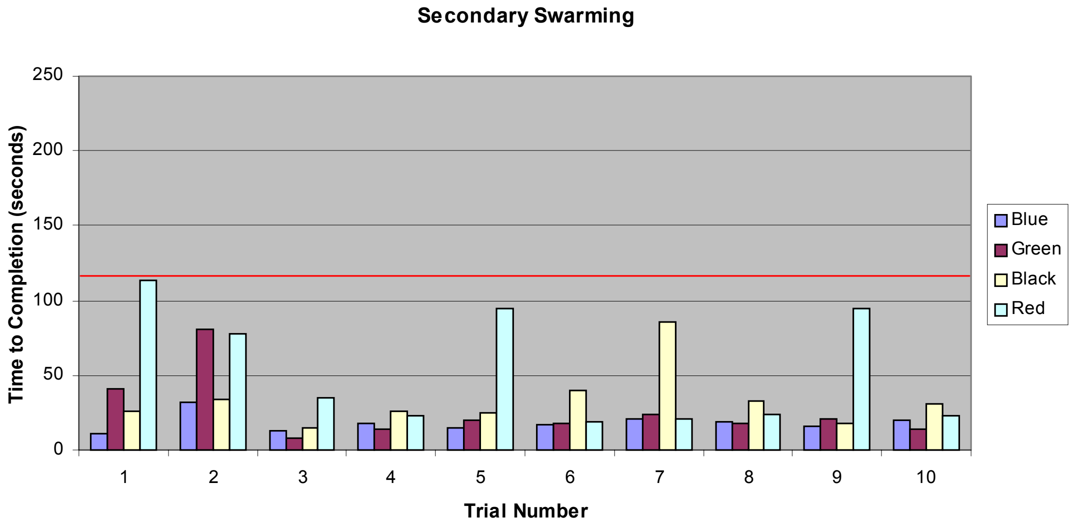

Figure 23. Secondary Swarming results

To show the secondary swarming process in action 12 images from the overhead camera have been used, these are shown in figure 24 frames A-L. Arrows have been added to the pictures to indicate what the blimps are homing in on. As shown in frame A all the blimps start in the same position as the previous experiment (section 4.2). When the beacon is switched on the green blimp acquires the signal and starts to home in (frame B), at the same time it starts emitting a secondary field from its rear facing IR transmitter. The nearby blimps that cannot detect the beacon soon lock onto the signal from the green blimp and follow it closer to their shared goal (frames C & D). In frame C the green blimp can be seen blocking (or occluding) the beacon signal preventing the blue blimp from homing in. However, this is no longer a problem as the secondary field emitted by the green blimp continues drawing the blue blimp closer, effectively making the green blimp invisible. When the green blimp eventually moves to one side the blue blimp is able to see the beacon and start homing in whilst simultaneously emitting its own secondary field. At this point (frames D&E) the red and black blimps are attracted to the rear of the blue blimp thus involving all the participants in the secondary swarming process. In frame F the black blimp acquires the beacon signal and starts emitting its secondary field, which in turn attracts the red blimp. The blue and black blimps continue homing in on the beacon dragging the red blimp closer as they do so (frames G & H). Finally in frame I the red blimp detects the beacon and quickly moves in to complete the experiment.

Frame A 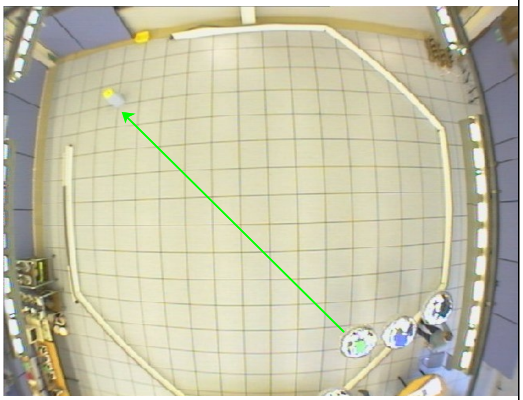	Frame B 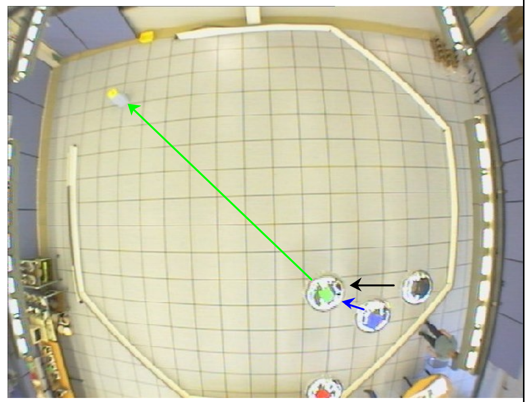
The green blimp acquires the beacon signal first and starts to transmit its secondary field and home in on the beacon.	The blue and black blimps immediately detect the secondary field from the green blimp and start to follow it toward the beacon.
Frame C 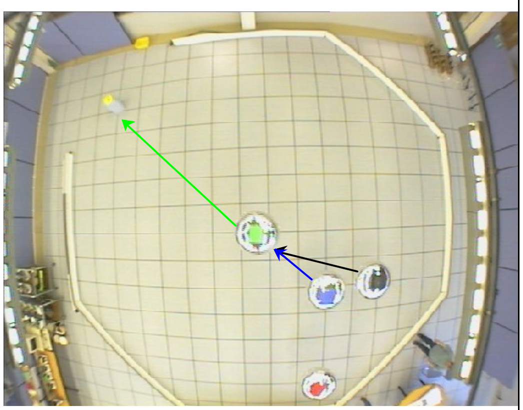	Frame D 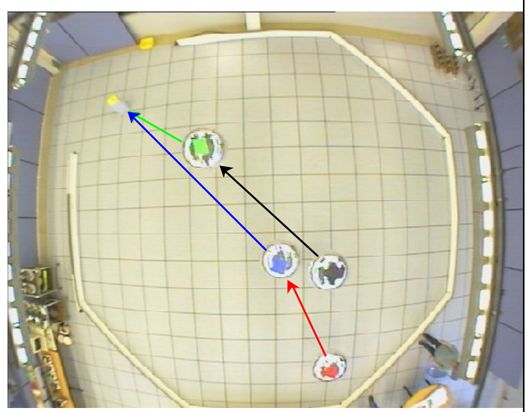
The red blimp unable to detect any signal remains floating around the start.	As the green blimp gets closer to the beacon it moves to one side giving the blue blimp a clear view of the beacon. At this point the red blimp detects the secondary field generated by the blue blimp.

<table>
<tr><td>Frame E
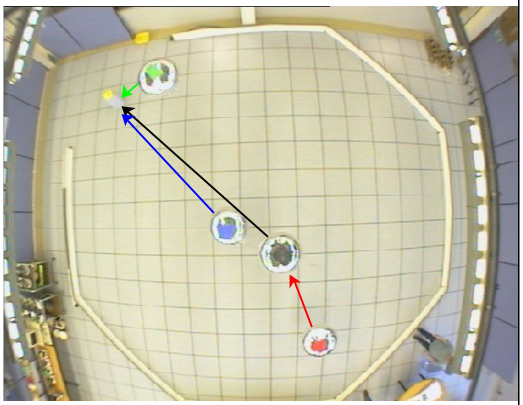</td><td>Frame F
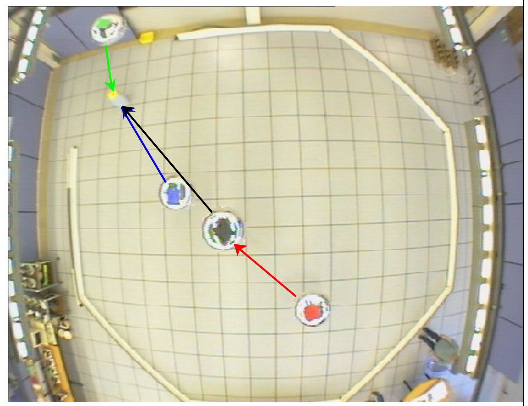</td></tr>
<tr><td>With the green blimp out of the way the black blimp is able to detect the beacon and start transmitting its secondary field.</td><td>The red blimp then starts to follow the stronger signal from the black blimp toward the beacon.</td></tr>
</table>

<table>
<tr><td>Frame G
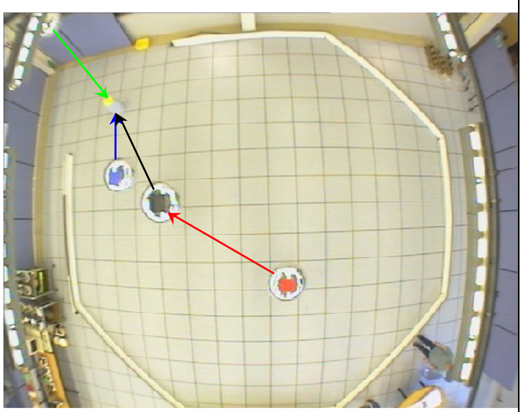</td><td>Frame H
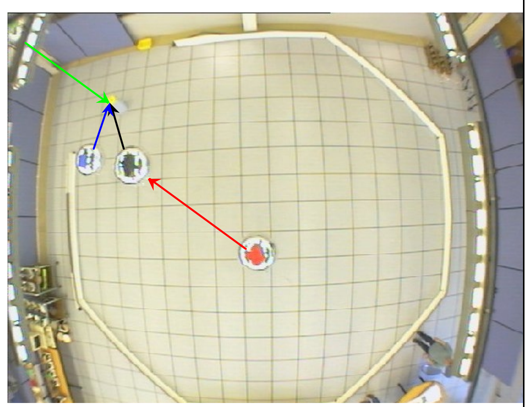</td></tr>
<tr><td>The blue and black blimps continue homing on the beacon drawing the red blimp ever closer.</td><td>The red blimp still following the black blimp moves closer to the beacon still unable to detect it.</td></tr>
</table>

Frame I 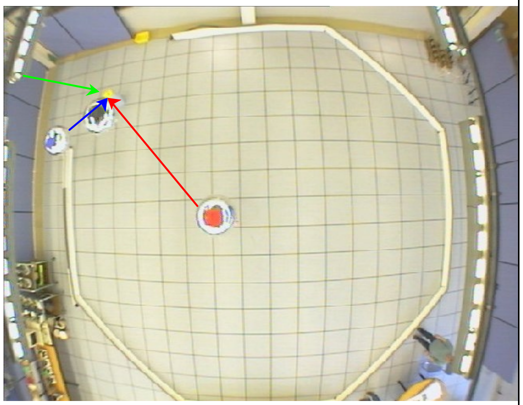	Frame J 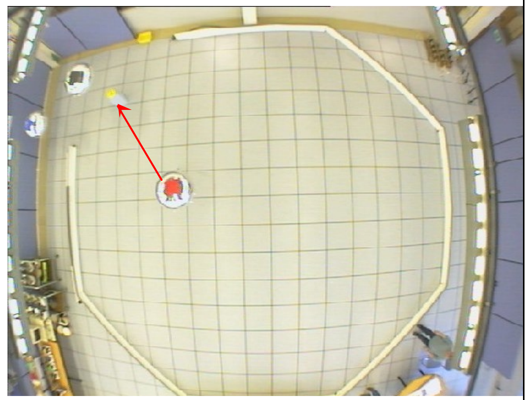
Finally with no other blimps obscuring the beacon from view the red blimp detects it.	The red blimp continues to home in while the others mill about on the other side of the beacon.
Frame K 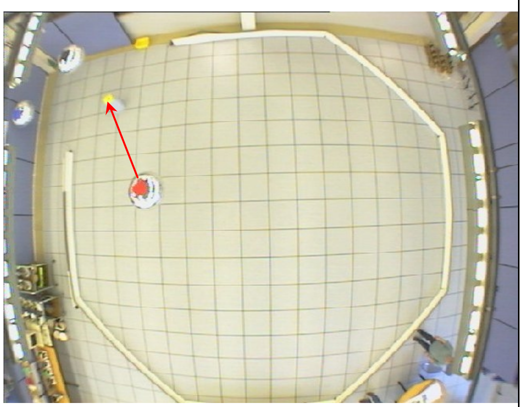	Frame L 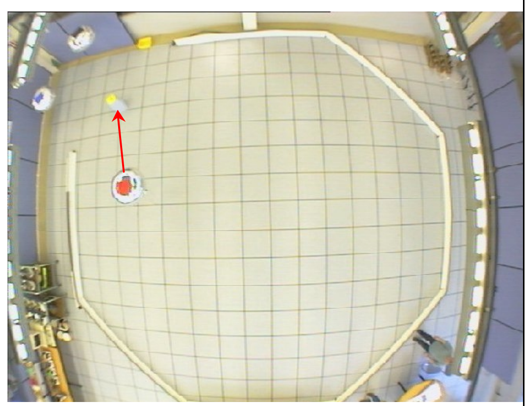
Slowly but surely the red blimp moves ever closer to its target.	Finally the last remaining blimp reaches its target

Figure 24. A sequence of frames showing Secondary Swarming

4.3.3 Conclusions

For all the trials that used the secondary swarming technique, 100% of the robots managed to manoeuvre to within two blimp diameters of the beacon within the given time period. In contrast using pseudoswarming a maximum of 75% of robots made it to the beacon within the time period. The difference in performance of the two algorithms is shown clearly in the graph (figure 25) below.

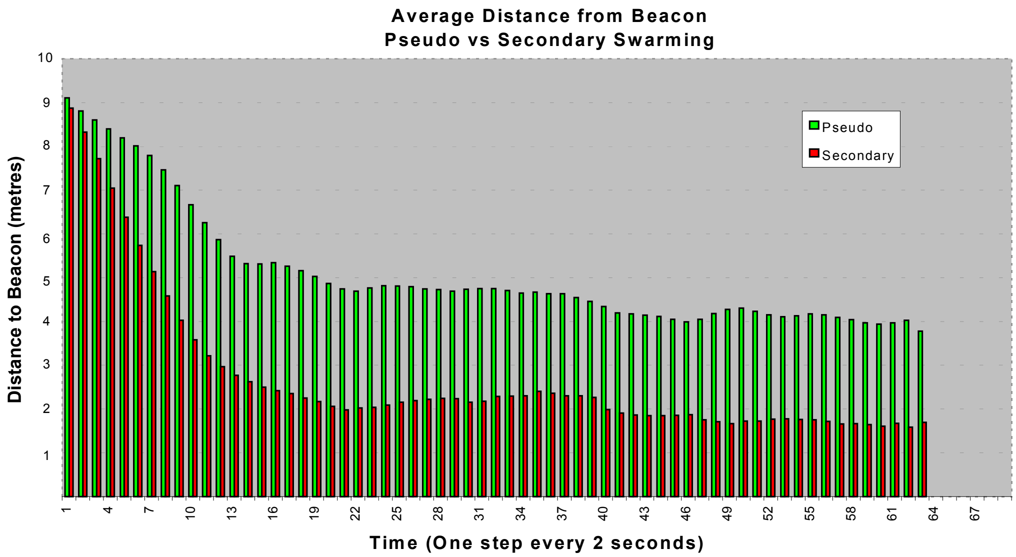

Figure 25. Showing average distance from the beacon for Secondary Swarming and Pseudoswarming

From the results (figure 22) it can be seen that the red blimp never reaches the beacon when pseudoswarming due to its sensor range being shorter than any of the other robots. This, however does not mean the robot is useless: by using the secondary swarming technique a less able robot can be guided towards the beacon by the other robots (as shown in frames D – L figure 24) and once in range it will detect the beacon and home in.

5. Discussion

The paper has shown that it is possible to implement the strategy of secondary swarming, initially studied in simulation, in real robots that move in a 3 dimensional environment and have 3 degrees of freedom. The blimp robots were chosen because they represented a challenging platform with a severely restricted payload. Minimalist strategies are therefore appealing since the weight restriction forced simple sensing, locomotion and communication techniques to be used.

The experiments showed that it is possible to significantly increase the performance of a group of 4 blimps when homing in on a static beacon starting in a low signal to noise domain. Two sets of experiments were conducted, the first using a purely individual homing strategy and the second using a collective strategy. Examination of the results showed that if a robot could not initially detect the beacon it had a very slim chance of finding it. However, using the collective strategy any robot detecting the beacon would transmit a secondary field effectively increasing the beacon's radius of influence, which would in turn attract robots outside the initial beacon range. The increase in performance between the two techniques is considerable since the use of a purely individual homing strategy (Pseudoswarming) led to no trials being completed within the time limit of 2 minutes. By comparison, every trial was completed within this limit using a collective strategy of secondary swarming.

6. References

Beckers R, Holland O & Deneubourg J-L. (1994). “From Local Actions to Global Tasks: Stigmergy and collective Robotics.” Proceedings of the Fourth International Conference on Artificial Life.

Feinleib M. (1980) "Photomotile Responses in Flagellates". In: Photoreception and Sensory Transduction in Aneural Organisms". Eds: Lenci F. & Colombetti G. pp 45-68. New York & London: Plenum.

Gulliford, S. (1999). “Emergent Group Behaviours of Multiple Autonomous Mobile Roots Using Minimal Sensing” MSc Dissertation, University of the West of England.

Holland, O.E. (1996) “ The use of social insect control methods with multiple mobile robots”. Control ’96, IEE, Exeter.

Holland,O.E. & Melhuish, C.R. (1996) “Getting the Most from the Least: Lessons for the Nanoscale from Minimal Mobile Agents” A-Life V, Nara, Japan 1996. Pages 59-66.

Kelly, I.D. & Keating, D.A. (1996) “Flocking by the fusion of sonar and active infrared sensors on physical autonomous mobile robots.” The third Int. Conf. on Mechatronics and Machine Vision in Practice. Guimaraes, Portugal. Vol. 1, pp 1/1 - 1/4.

Kube, C & Zhang, H. (1992). “collective Robotic Intelligence” Proceedings from the Second International Conference on the Simulation of Adaptive Behaviour MIT Press.

Mataric, M.J (1994). “Interaction and Intelligent Behaviour” PhD Thesis. MIT.

Mataric, M.J. (1992). “Designing Emergent Behaviours: From Local Interactions to Collective Intelligence.” Animals to Animats 2.

Melhuish, C.R & Holland,O.E. (1996) “Getting the Most from the Least: Lessons for the Nanoscale from Minimal Mobile Agents” Artificial Life V. Proc. 5th Int. Workshop on the Synthesis and Simulation of Artificial Living Systems, Nara, Japan 1996.

Melhuish C. (1999) “Employing Secondary Swarming with Small Scale Robots: A biologically Inspired Collective Approach”. in Proceeding of 2nd Internatational Conference on Climbing and Walking Robots – CLAWAR. Portsmouth U.K. Professional Engineering Publishing Ltd.

Melhuish C. (2001) “Strategies for Collective Minimalist Mobile Robots” Engineering Research Series. Professional Engineering Publishing. ISBN 1 86058 318 0.

Melhuish C., Holland O. and Hoddell S. (1998a) “Using Chorusing for the Formation of Travelling Groups of Minimal Agents” 5th International Conference on Intelligent Autonomous Systems, Sapporo,Japan.

Melhuish, C., Holland, O. and Hoddell, S. (1998b) “Collective Sorting and Segregation in robots with Minimal Sensing” 5 th International Conference on the Simulation of Adaptive Behavior. Zurich. From Animals to Animats. MIT Press.

Reynolds, C.W. (1987) “Flocks, herds and schools : A distributed behavioural model.” Computer Graphics : Proceedings of SIGGRAPH '87. Vol. 21, no. 4.