Robot Sheepdog Project achieves automatic flock control

Richard Vaughan*, Neil Sumpter, Andy Frost***, Stephen Cameron***

* Oxford University Computing Laboratory,Parks Road, Oxford, UK

Richard.Vaughan@comlab.ox.ac.uk Stephen.Cameron@comlab.ox.ac.uk

** School of Computer Science, University of Leeds, Leeds, UK

neils@scs.leeds.ac.uk

*** Silsoe Research Institute, Wrest Park, Silsoe, Bedfordshire, UK

Andy.Frost@bbsrc.ac.uk

Abstract

This paper describes a mobile robot that can enter a circular arena, gather a flock of ducks and manouvre them safely to a specified goal position. A minimal simulation model of the ducks' flocking behaviour was developed and used as a tool to guide the design of a general flock-control algorithm. The algorithm was first tested in simulation then tranferred unchanged to a physical robot which succeeds in gathering a real flock of ducks. This is the first example of a robot system that exploits and controls an animal's behaviour to achieve a useful task.

Robots successfully manipulate many objects in factories and laboratories. Research continues on manipulating objects with complex or variable shapes and dynamics eg. food products (Juste et al., 1997). Recent work in mobile robotics has focused on 'adaptive behaviour' in animals, in order to extend the abilities of robots (Maes, 1990) (Hallam and Hayes, 1994), and to better understand the processes occurring in real creatures (Webb, 1994)
(McFarland and Bosser, 1993). The Robot Sheepdog Project examines the robotic manipulation of animals by exploiting their adaptive behaviour. In contrast, previous work combining robots and animals (Trevelyan's robot sheep-shearer (Trevelyan, 1992), Silsoe Research Institute's milking robot (Frost et al., 1993)) has deliberately minimised animal behaviour by physical restraint. We have demonstrated a mobile robot that can enter a circular arena, gather a flock of ducks and manouvre them safely to a specified goal position. This is the first example of a robot system that exploits and controls an animal's behaviour to achieve a useful task.

The sheepdog's gather-and-fetch task was chosen because of its familiarity and the strong interaction between the dog, shepherd and flock animals. Using ducks instead of sheep allows us to experiment on a conveniently small scale, in a controlled indoor environment. Duck flocking behaviour is recognised by shepherds as similar to sheep; ducks are often used to train sheepdogs because of their relatively slow movement.

Figure 1 Sheepdog with duck flock in Lancashire, 1996.

Flocking is considered an adaptive behaviour, as it affords various advantages in hazard-avoidance, mating and foraging. Models of flocking behaviour exist in the literature and are generally derived from Hamilton's observation that flocking may be produced by the mass action of individual animals, each seeking the proximity of its nearest neighbours (Hamilton, 1971). It was later suggested that this behaviour can be well modelled by an attractive 'force' acting between the animals, with the magnitude of the attraction varying with the inverse square of the animals' mutual distance (Partridge, 1982) (Warburton and Lazarus, 1991). It is argued that this relationship represents a linear response to sensory information which itself varies with the inverse square of distance. Similar models have produced realistic computer animations of bird flocks (Reynolds, 1987).

These ideas are familiar in robotics, where such *po-*

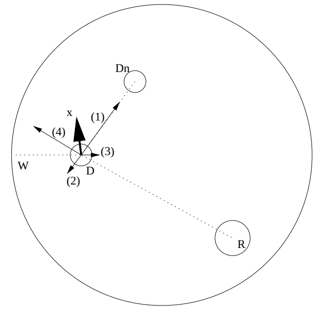

$$\mathbf{d} = \sum_{n=1}^{N} \left(\left(\frac{K_1}{|\mathbf{DD_n}|^2} \right) \widehat{\mathbf{DD_n}} - L - \left(\frac{K_2}{|\mathbf{DD_n}|^2} \right) \widehat{\mathbf{DD_n}} \right) - \left(\frac{K_3}{|\mathbf{DW}|^2} \right) \widehat{\mathbf{DW}} - \left(\frac{K_4}{|\mathbf{DR}|^2} \right) \widehat{\mathbf{DR}}$$

(1) (2) (3) (4)

Figure 2 Flock model (schematic not drawn to scale). Key: gain parameters $K_{1\to 4}$; offset parameter L; number of ducklets N; ducklet position D, other ducklet D_n; Robot position R; Nearest point on wall W; algorithm terms $(1 \to 4)$ and resultant $\mathbf{d}$ (where $\widehat{\mathbf{a}}$ is the unit vector of $\mathbf{a}$).

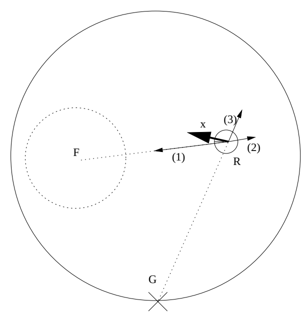

$$\mathbf{r} = K_1 \mathbf{RF} - \left(\frac{K_2}{|\mathbf{RF}|^2} \right) \widehat{\mathbf{RF}} - K_3 \widehat{\mathbf{RG}}$$

(1) (2) (3)

Figure 3 Flock control method (schematic not drawn to scale). Key: gain parameters $K_{1\to 3}$; flock center F; Robot position R; algorithm terms $(1 \to 3)$ and resultant $\mathbf{r}$ (where $\widehat{\mathbf{a}}$ is the unit vector of $\mathbf{a}$)

tential field techniques are used for navigation (Cameron and Probert, 1994). This class of algorithm uses the analogy of forces acting on particles, such that the robot will move as if it were a particle attracted or repelled from features in its environment. A robot is typically attracted to a goal position and repelled from obstacles. The commonality of these animal and robot behaviour models forms the basis of an effective flock-gathering strategy.

A minimal simulation model of the duck-herding scenario was created, in which six model ducks (ducklets) move in a circular arena containing a model robot. A ducklet's movement vector $\mathbf{d} = (\frac{\delta x(t)}{\delta t}, \frac{\delta y(t)}{\delta t})$ is determined by the function shown in Figure 2.

The ducklets are (Figure 2, term 1) attracted to each other, aggregating the flock; (2) repelled from each other, preventing collisions and maintaining inter-ducklet spacing; (3) repelled from the arena wall, preventing collisions. A further term (4) which produces repulsion from the robot is proposed to model the aversive response of the ducklets to the robot. All these forces are scaled according to the inverse square of distance, and each ducklet moves according to the resultant of the forces acting upon it. The simulation produces a realistic-looking flock which can be manipulated by steering the model robot.

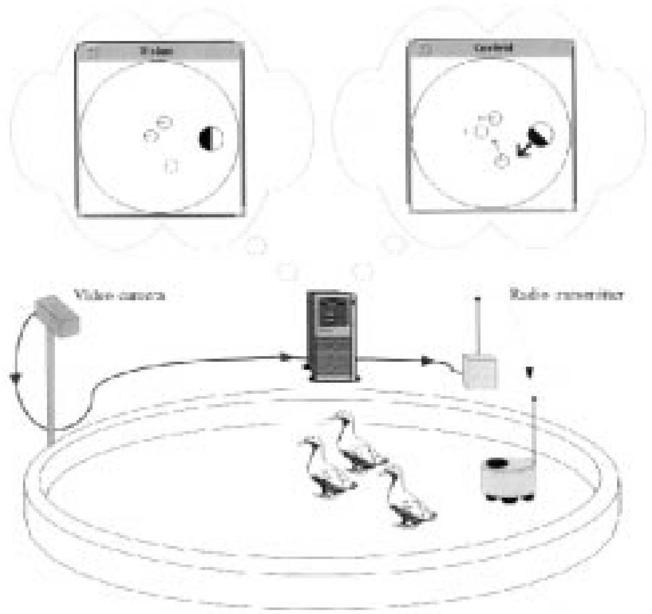

Figure 4 Robot Sheepdog system overview.

By experimenting with the simulator, a closely related algorithm was developed which steers the robot to gather the flock and return it to a goal position on the edge of the arena. The robot's movement vector **r** is given by the function shown in Figure 3. The robot is (Figure 3, term 1) attracted to the center of the flock (defined as the average position of all the ducklets) with a magnitude proportional to their mutual distance. This force causes the robot to move towards the flock. A second force (2) repels the robot from the flock center with a magnitude proportional to the inverse square of their mutual distance. This prevents collisions. The resultant of these two forces creates a circular orbit of zero potential around the flock centre. A further force (3) repels the robot from the goal position with a constant magnitude. This has the effect of tilting the potential landscape such that the orbit around the flock now has a minimum behind the flock with respect to the goal. The robot will move towards this point; driving the flock away from it and towards the goal. A separate mechanism is used to prevent the robot hitting the walls (not described here).

This method was tested using the flock simulation. A representative experiment is shown in Figure 6: The robot starts at the goal position, approaches the ducks, moves around behind them with respect to the goal, driving them in the desired direction. The system stabilises with the ducklets near the goal, and the robot standing off some distance away. The paths of the robot and flock center are shown (Figure 6 top) and the distance of the flock from the goal (our success metric) is plotted over time (Figure 6 bottom).

The method was found to be successful and robust over a wide range of flock parameters (Vaughan et al., 1997). The algorithm was then implemented on a physical robot in order to test its performance with a real flock of ducks.

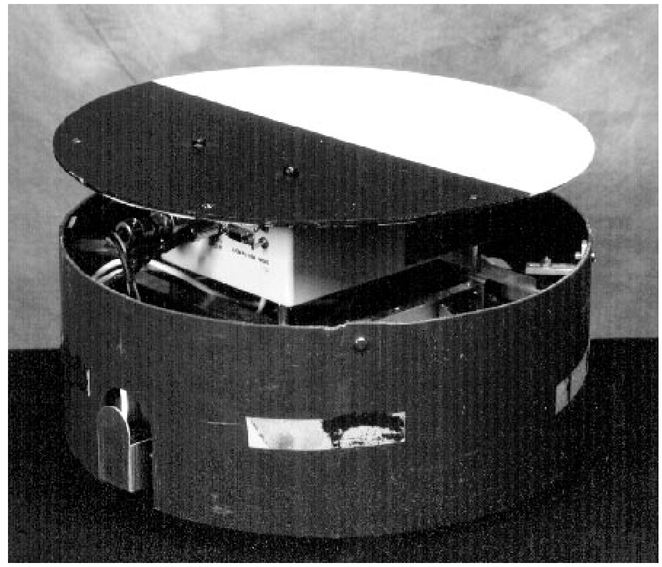

Figure 5 'Rover', the RSP vehicle.

The physical experimental system comprises a robot vehicle, a workstation and a video camera (Figure 4). The vehicle is designed to work in a duck's environment: outdoors, on short grass, and in real time. Thus our robot has a top speed $\approx 4ms^{-1}$ (roughly twice as fast as our ducks) and acceleration $\approx 1ms^{-2}$. It is covered in a soft plastic bumper mounted on rubber springs, ensuring duck safety. In the tradition of mobile robotics, we call it 'Rover' (Figure 5).

The vehicle and the ducks are free to move in a visually uniform arena of 7m diameter, in view of the overhead camera. The position and orientation of the robot, and the position and size of the flock are determined by processing the video image stream . A standard background-differencing technique is used, to achieve an update rate of approximately 25Hz. The vision system's output is used to generate a desired vehicle trajectory using the algorithm described above.

The vehicle's current speed and heading are compared to the desired movement vector **r**, and new wheel-speeds (R_{left}, R_{right}) are determined by the functions

$$\begin{pmatrix} R_{left} \\ R_{right} \end{pmatrix} = \begin{pmatrix} |\mathbf{r}| + \frac{D}{2}(\theta - \angle\mathbf{r}) \\ |\mathbf{r}| - \frac{D}{2}(\theta - \angle\mathbf{r}) \end{pmatrix}$$

where θ is the robot's current heading and D is the distance between the wheels. These wheel-speed demands are passed to the vehicle via a radio modem. The proportional controller

$$u(t) = K(R(t) - E(t))$$

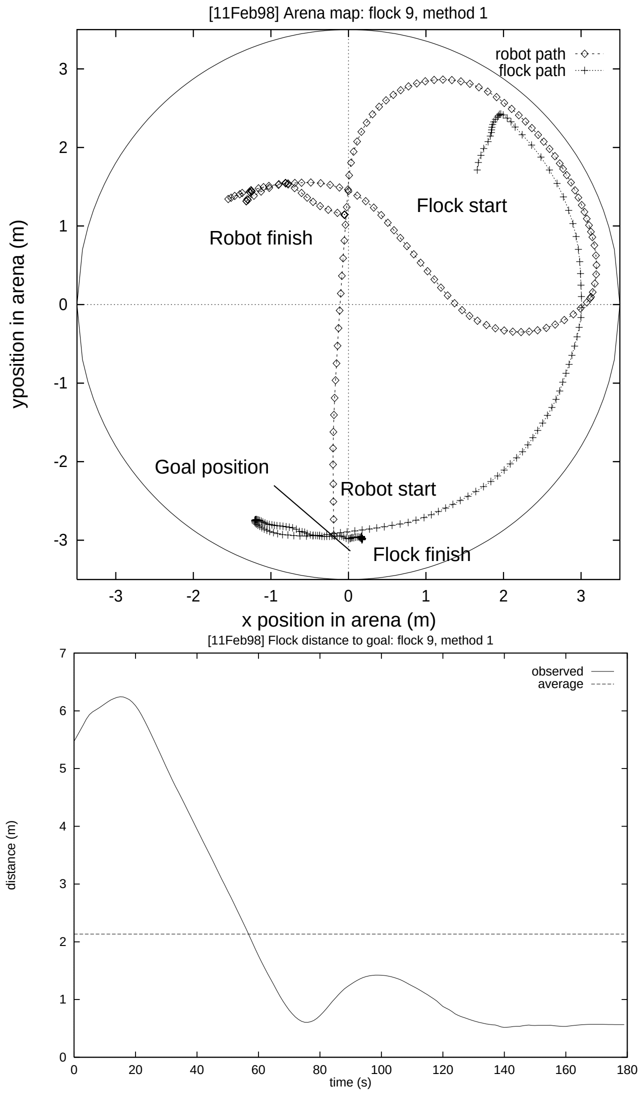

Figure 6 Simulation results. Top: paths of vehicle and flock (goal is at (0, -3.3). Bottom: distance from flock centre to goal position over time.

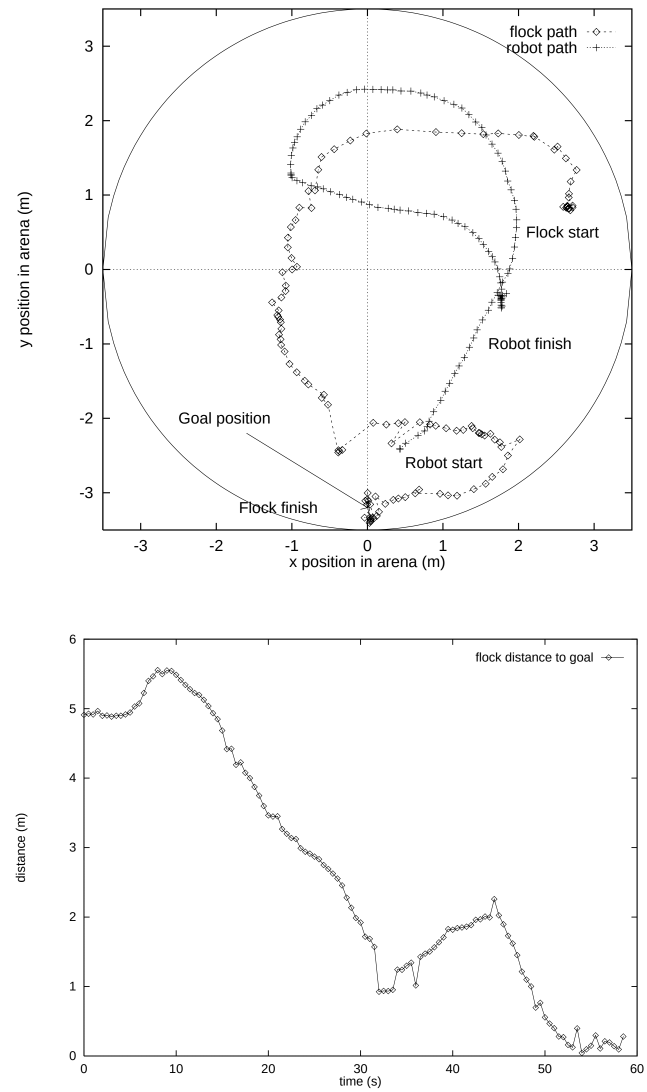

Figure 7 Real-world results. Top: paths of vehicle and flock (goal is at (0, -3.3). Bottom: distance from flock centre to goal position over time.

(where u is the output to the wheel, R is the desired speed, E is the speed error, and K is the controller gain) is implemented by the vehicle's on-board systems for each wheel at 100Hz. The controller gain was chosen by experiment and it is found that the vehicle's movement closely approximates the desired trajectory.

A random point along the arena boundary was chosen as the flock goal, corresponding to position (0, -3.3) on Figure 7 (top). With the robot inactive and positioned near the flock goal, the ducks were introduced into the arena and allowed to move freely. After 5 minutes accommodation time, the robot was activated. Figure 7 (top) shows the paths of the robot and the flock center over a 60 second trial, with samples taken at 0.5s intervals. Figure 7 (bottom) shows the distance of the flock center from the goal point over time; it can be seen that the flock is effectively moved to the goal position. The behaviour of the real-world system is subjectively similar to the simulation.

Since the first draft of this paper this pilot trial has been followed by extensive experimental trials, and a superior algorithm has been developed and tested. Papers describing this recent work will appear during 1998.

In this paper we have described a robot system that achieves a sheepdog-like task, gathering and fetching live animals to a pre-defined goal position. We believe this is the first automatic system to exploit an animals adaptive behaviour to achieve a useful task. The robot's controller was designed and tested using a minimal simulation model of the ducks' flocking behaviour, and successfully transferred to the real world. We conclude that behavioural simulations can be plausible engineering design tools, and that such a methodology is appropriate for such animal-interactive robotics experiments. These results also support the attraction-force model as a good description of flocking behaviour.

References

Cameron, S. and Probert, P. (1994). *Advanced Guided Vehicles: Aspects of the Oxford AGV project.* World Scientific.

Frost, A., Street, M., and Hall, R. (1993). The development of a pneumatic robot for attaching a milking machine to a cow. *Mechatronics*, 3(3):409–418.

Hallam, B. and Hayes, G. (1994). Comparing robot and animal behaviour. Research Paper 598, Department of Atrificial Intelligence, University of Edinburgh.

Hamilton, W. (1971). Geometry for the selfish herd. *Journal of Theoretical Biology*, 31:295–311.

Juste, F., Andreu, G., Valiente, J., and Benlloch, J., editors (1997). *Proceedings of the International Workshop on Robotics and automated machinery for bio-productions.*

Maes, P., editor (1990). *Designing Autonomous Agents.* The MIT Press: Cambridge, MA.

McFarland, D. and Bosser, T. (1993). *Intelligent behavior in animals and robots.* MIT Press.

Partridge, B. L. (1982). The structure and function of fish schools. *Scientific American*, pages 114–123.

Reynolds, C. W. (1987). Flocks, herds and schools: A distributed behavioural model. *Computer Graphics*, 21(4):25–34.

Trevelyan, J. P. (1992). *Robots for Shearing Sheep: Shear Magic.* Oxford University Press.

Vaughan, R., Henderson, J., and Sumpter, N. (1997). Introducing the robot sheepdog project. In *Proceedings of the International Workshop on Robotics and Automated Machinery for BioProductions.*

Warburton, K. and Lazarus, J. (1991). Tendency-distance models of social cohesion in animal groups. *Journal Of Theoretical Biology*, 150:473–488.

Webb, B. (1994). Robotic experiments in cricket phonotaxis. In Cliff, D., Husbands, P., Meyer, J., and Wilson, S., editors, *From Animals to Animats 3: Proceedings of the third international conference on the simulation of adaptive behaviour*, pages 45–54, Cambridge, Mass. MIT Press.