Simulating Swarm Behaviour of Robots

K.A.Hawick and H.A.James
Computer Science Division, School of Informatics
University of Wales, Bangor, North Wales, LL57 1UT, UK
{hawick,heath}@bangor.ac.uk
Tel: +44 1248 38 2717, Fax: +44 1248 36 1429

ABSTRACT

We describe simple simulation models of collective swarm behaviour of simple robots. We have experimented with mixed avoidance and seeking behaviour using tri-wheeled robots using simple ultrasonic proximity sensors. Our model assumes a radial distance of attraction or repulsion and supports multiple species in the swarm based on sensor recognition protocols. We describe the surprising success of our model in predicting the gross and statistical behaviour of our swarm and discuss limitations and potential improvements to the model. We also discuss scaling experiments to large swarm sizes. We also speculate on thermal physics model analogies to the swarm and how these can be parameterised.

KEY WORDS
robots and automation; swarm behaviour; simulation model

1 Introduction

The problem of understanding collective and emergent behaviour in swarms of robots is a fascinating one and has recently become more accessible to research work due to commodity priced robot parts and embedded control electronics. We are investigating cooperative behaviour amongst a swarm of small tri-wheeled robots based on Warwick's Cybot mechatronics[8], but with Tini-based control processors[3] running Java software[5].

Figure 1 shows a photograph of a typical Cybot robot. The device itself is around 30cm in length and a simplified schematic of our modified robot is shown in figure 2.

Figure 1. Real Robots "Cybot" Mechatronics showing the ultrasonic proximity sensors.

In this paper we describe some simulations of collective behaviour amongst these simple systems. Our goal is to develop a valid simulation model for collective avoidance and following behaviour that can be parameterised by proximity range settings. We are particularly interested in the statistical and emergent properties of groups or swarms of robots. We describe our mechatronics platform and its features in section 2 and the proximity sensors and our modifications to them in section 3. We outline the programmable behaviour of our swarm robot in section 4.

We present our simulation model in section 5 and some extensions to model segregation behaviour in section 6. One experiment we are using our simulation model to investigate is that of grouping or phase separation behaviour. We consider two types of robot. All robots might have identical mechatronics and electronics but are marked in some way so that for example intra group members attract or seek out each other while inter group members repel or avoid one another.

We discuss some quantitative metrics for comparung predicted behaviour with that or the robot swarm. One metric we find useful is that of the peak or mean value in the structure function. We compute this as a Fourier transform of the spatial correlation function. Finally we review the limitations of our model in section 7.

2 Robot Mechatronics

The Cybot toy comes equipped with motors and PIC microcontroller[7] electronics allowing several simple behaviours to be readily programmed. The PIC series of controllers are relatively cheap and can by programmed in assembler code or in C. We are interested in programming different high level behaviours into a swarm of cybots and have developed a more sophisticated control architecture that uses the Tiny InterNet Interface (Tini) controller developed by Dallas Semiconductors. This system can be programmed in Java and is not only able to communicate with the PIC controllers we use for our low level device drivers,

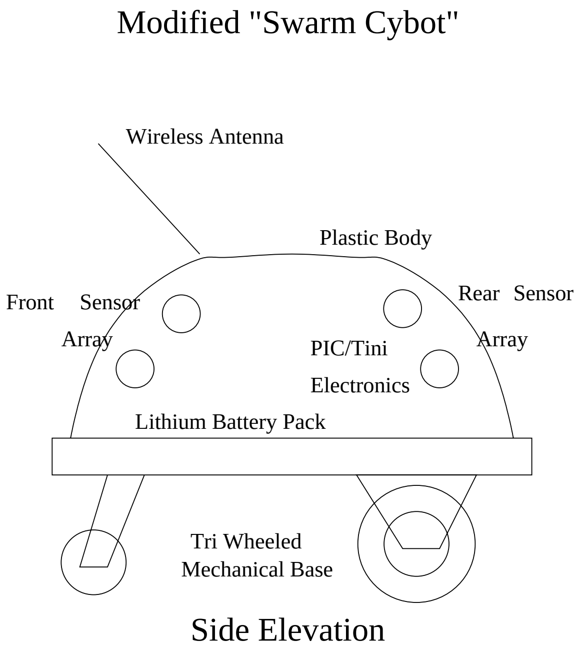

Figure 2. Schematic structure of our modified Swarm Cybot using PIC microcontroller and Tini electronics running the DISCWorld Lite Daemon on the Java Virtual Machine on the Tini board..

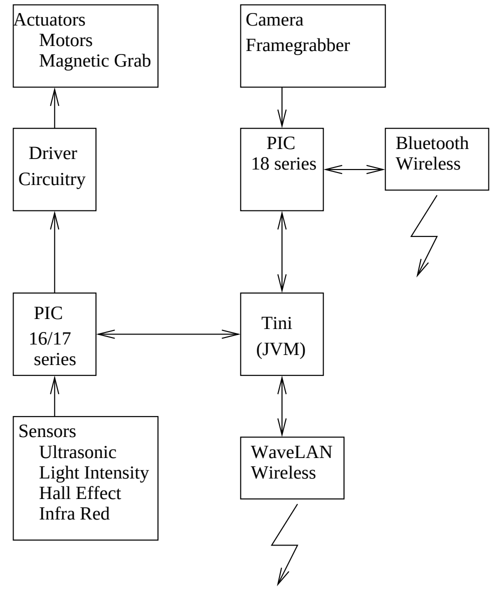

Figure 3. Control system block structure for modified Swarm cybot node.

Figure 4. Cybot Pen with WebCam "positioning satellite" above.

but also has ethernet capability.

Figure 3 shows the block architecture of our swarm node robot. We employ the cheaper series 16 and 17 series PIC chips to drive our motors and magnetic actuators. We are experimenting with the more sophisticated (in terms of memory and interfaces) series 18 PIC chip to interface to a frame-grabber and camera assembly. Although there are several good camera systems available we require one that is both light enough and cheap enough that we can employ one on all our swarm robot nodes. It may be however that we design special purpose "seeing class" nodes that are camera enabled.

Our full architecture as shown in figure 3 is still in development, but not all of it is needed for the experiments reported in this paper. We are developing more sophisticated control electronics for the Cybot robots using Tini control boards and Bluetooth wireless radio[2] communications. We have developed a lightweight Java based messaging system[6] that will allow simple inter robot communications strategies to be programmed.

Figure 4 shows a custom rectilinear enclosure built for our modified cybots to interact. The enclosure allows up to 8 cybots to interact closely with their ultrasonic sensor sensitivities turned down to around 50cm range or less. There are various thermal physics analogy games to play with such an enclosure including "suddenly" removing one wall and allowing the robots to do a "free expansion" into a greater volume.

A key feature of the enclosure is that a camera is mounted overhead and can readily be used to monitor and track individual robots as they move around. We experimented with various colours and backgrounds for the enclosure. The photograph shows white shelving units and a base of white linoleum. This does give good contrast with the robots but requires a filter to avoid camera signal saturation.

3 Sensors

We have carried out considerable experimentation to determine just what sensors the individual robots need carry. Since our goals are to investigate collective behaviour amongst a medium to large number of robots, it is particularly important to minimise the individual cost of a node. It is also important to use commodity parts to ensure similarity amongst nodes. We want any differences to be due to programmed or learned behaviour and not to the accidental physical characteristics that arise from manufacture.

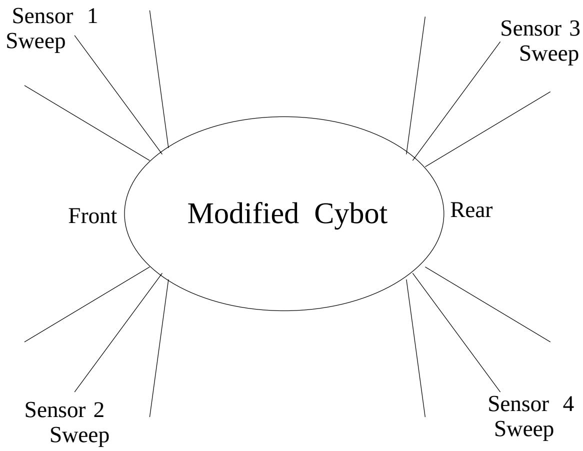

Figure 5. Sensors around the Modified Cybot base

Figure 5 shows the arrangement of ultrasonic sensors around our modified cybots. The standard Cybot comes with two forward sensing pairs of ultrasonic transceivers that can be tuned in sensitivity with analogue electronics to report a high/low proximity signal to the robots control processor. These proximity alarm signals can be used to determine robot behaviour. The forward facing sensors are adequate for collision avoidance of a single robot. One of the weaknesses of this model is the lack of awareness of robots approaching from behind.

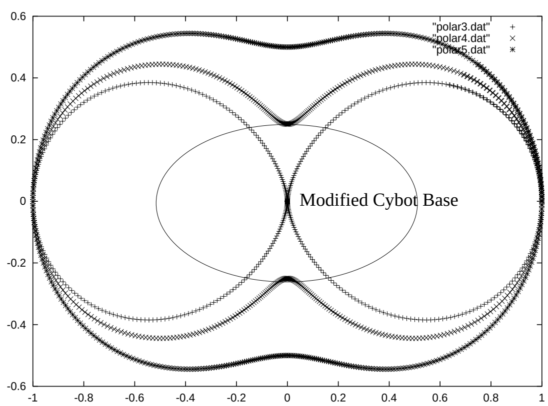

Figure 6. Contours of proximity range around the modified cybot

Figure 6 shows the contours of proximity sensitivity around our modified cybot. Simple intensity measurements sufficed to verify that this can be adequately modelled using a $\cos^2(\theta)$ function, with front and rear of the cybot assumed symmetrical. This sensor arrangement allows us to consider forward and rear avoidance motion as symmetric. It is possible to incorporate the elliptical asymmetry into our model but it is not clear if this is necessary. The simplifying model of circular robots is likely adequate except for very dense robot packing experiments.

This all round proximity sensor is important if we are to treat robots using the model described in section 5. The robot is now able to take responsibility for rear collisions a s well as forward facing ones. Furthermore it has enough information to correct its path during a turn.

4 Robot Behaviours

The modified cybot supports basic movement control and enhanced proximity detection. Proximity signals are updated frequently enough to allow differentiation of a static obstacle or a moving one. The behaviours we program into our robots for the investigations described in this paper can be summarised as:

```
move  forwards at medium  speed until
  static proximity  alert occurs (indicating the enclosure)
    in which  case  choose turning angle  and change direction
  dynamic proximity  alarm occurs(ie another  robot)
    in  which case turn towards it unless it is  too close
       in which case  reverse (collide)
```

A variation discussed in section 6 is to determine which species or tribe a dynamic target is and avoid it completely (regardless of its distance) if it is of another tribe.

These high level descriptions translate into a simple process model to control the two drive wheels of the robot.

5 Simulation Model

Our goal is to model medium to large configurations of robot systems using behaviour models derived from simple physics. One starting point is to consider the robots randomly moving around in their enclosure as analogous to "two dimensional molecules" interacting. The simplest model is to treat the robots as essentially random, and their motion to be Brownian. Qualitatively a number of robots set in the enclosure do behave this way, but it is more interesting to determine a quantitative metric for their collective behaviour.

The robots can be programmed to carry out simple collision avoidance but to seek out a target that is in motion, using simple time difference calculations from sensor signals. This behaviour is roughly analogous to a set of molecules that are weakly attractive to one another and with might be expected to display clumping behaviour.

We attempt to describe this behaviour using a model Hamiltonian was constructed using the Lennard-Jones form [1]:

$$U(R) = 4\epsilon \left[\left(\frac{\sigma}{R} \right)^{12} - x^{i,j} \left(\frac{\sigma}{R} \right)^{6} \right] \qquad (1)$$

This gives the potential energy U due to the pair-wise interaction of two robots of species i, j, separated by some radial distance $R = |\mathbf{r}_i - \mathbf{r}_j|$, given two robot specific parameters in the form of an energy ϵ and a length scale σ.

We have also incorporated a cross term coupling fraction $x^{i,j}$ which can be set greater for interactions between like species i, j than for opposite species. This segregation model is described in section 6.

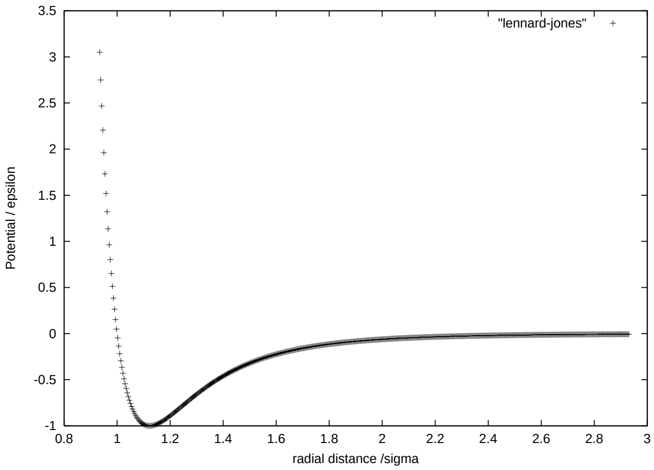

Figure 7. Lennard Jones Potential Model showing hard core exclusion zone and weakly attractive long distance tail.

Figure 7 shows how the potential energy varies with distance, between a pair of "molecules" or robots using the simple Lennard-Jones potential. A hard core is implemented by having the potential rise to infinity at zero separation with a steep curve to model the fact that two robots cannot physically overlap. A weakly attractive potential tail models the behaviour that robots should cluster together, and a potential minima models the preferred mean separation distance.

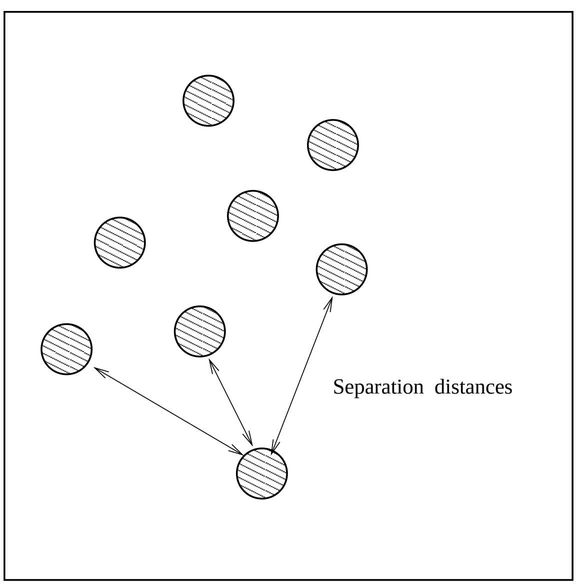

Figure 8. Eight Robots simulated in Pen, and tracked with camera.

Figure 8 shows a typical configuration of eight robots snapshot-ted in the enclosing pen. A clustering tendency can be observed once the robots have "thermalised" in the pen. This can be investigated quantitatively by tracking the central positions of robots using simple pattern recognition on samples from the overhead camera. The correlation function $c(r)$ is histogramnmed from the distribution of separation distances between all pairs of robot nodes. This can then be Fourier transformed to obtain a structure function $S(\mathbf{q})$. This is expressed by:

$$S(\mathbf{q}, t) = \frac{1}{N} \sum_{\mathbf{r}} e^{i\mathbf{q}.\mathbf{r}} \sum_{i} \langle [c_i(\mathbf{r_i}, t) - \overline{c}][c_i(\mathbf{r_i} + \mathbf{r}, t) - \overline{c}] \rangle \quad (2)$$

where q is the wavevector in reciprocal space and is given by $q = \frac{2\pi}{r}$ if r is a distance in the position space of the configuration. The c_i are the concentration variables and $\overline{c} =< c >$ is the global mean concentration and is equal to 0.5 for a 50% system. $S(\mathbf{q}, t)$ is effectively a structure function characterising the domains in the system. It is most useful if it is averaged spherically to give $S(q, t)$.

This sort of numerical scattering experiment can be carried out in condensed matter simulations and is a useful means of obtaining a characteristic length scale from a cluster. In this size of system the method is rather crude, but a definite peak or maximum is observed in the structure function as the seeking robots thermalise. The peak found corresponds to an approximate mean separation of 50cm, which is consistent with observed behaviour. The analogue electronics sensitivity of the ultrasonic proximity sensors is not easy to tune with any degree of accuracy, but the gross behaviour is correct in that lowering the sensitivity means the robots drift further apart and heightening it means then can cluster more closely. This is shown in changes in the measured peaks in the structure function. We plan to repeat this experiment with a larger number of robots.

6 Segregation Behaviour

The model can be made more sophisticated if there are more than one species of molecule or robot present. Entities of the same species can use the model shown whereas entities of dissimilar species can use a simple exponential repulsive potential, without the attractive part of the Lennard-Jones model. This multi-species model can be used to model phase separation behaviour[4].

Figure 9 shows results from a 2-d disk based molecular dynamics simulation of a swarm of 64 robots, initialised randomly and allowed to interact for some 10^7 integration time steps, using the Beeman integration algorithm [1]. This corresponds approximately to 5 minutes of real time for our robots which travel at rather less than 0.5 meter per second.

This model is rather appealing as it reproduces the gross behaviour of the robot swarm despite a number of major simplifications. This model was adapted from a similar one used to study phase separation behaviour amongst

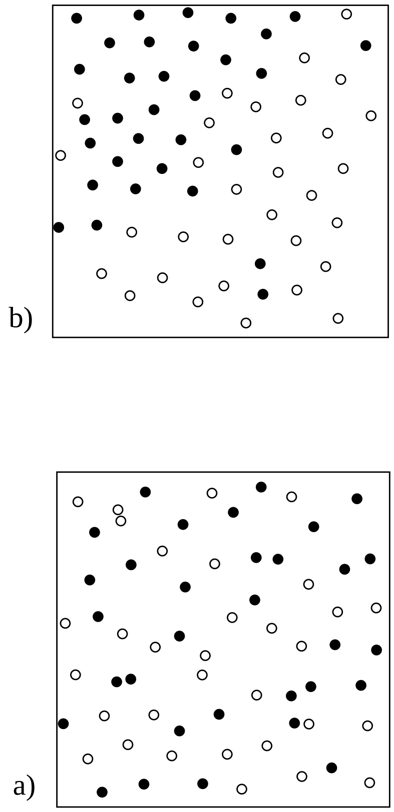

Figure 9. Simple 2d disk simulation using Molecular Dynamics approach and 2d Lennard-Jones potential model with alternating signs for black/white units.

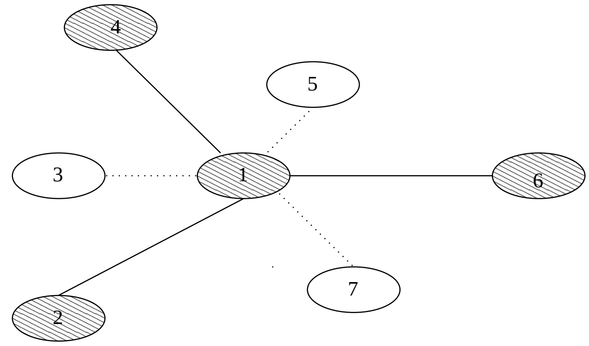

Figure 10. Location/Identification Graph for node 1. Two groups or types of robot are shown - one shaded, one white.

atomic species in a quenched metal alloy[4]. It is not yet clear to us how well the actual time and length scales can be made to model the robot swarm nor how closely simulation model parameters need match that of individual robots.

To investigate this behaviour in our robot swarm we need individual robots to recognise those in proximity to them and to adjust their behaviour accordingly. One mechanism we are exploring is for individual robots to carry a transponder such as a coloured light or some other recognisable token to denote which "tribe" they belong to. This requires development of further on-board custom electronics. In the interim we are experimenting with a system that is rather more general. Robots are equipped with wireless communications and are given accurate location information and a description of their neighbours.

This approach is not unrealistic as real world robots might well have global positioning systems inbuilt in the near future. This is certainly a quick fix to our problem of identifying nearby robots.

Figure 10 shows the identification and location graph for the nodes in proximity to a particular robot (numbered 1). The robot has to essentially maintain a world map of those other robots in its immediate vicinity. The diagram shows a case where two tribes of robots are interspersed and robot number 1 must decide how to react to different neighbours. It is assumed it can recognise nearby robots as belonging to the two different groupings and even if they are too close to react to individually it might be feasible for the robot to move towards a perceived mass of its own tribe a nd away from a mass of the enemy.

Various sensor combinations could implement this identification of which group a nearby robot is in. We plan to investigate more localised sensor options for this but at present we are able to utilise overhead camera (pseudo satellite) information as well as simple pattern recognition of a coloured disk mounted on top of the robots to identify them. This information can be broadcast to all robots over the wireless link. Consequently they can in fact build up a very accurate and complete world map of their neighbours and act accordingly.

At present our system is able to broadcast a complete world map update in around 1/10th of a second. This is mostly limited by the recognition software at present, although with a very large number of participating robots, the broadcast communications is likely to dominate.

7 Discussion and Conclusions

The results of our simple model are encouraging and we are working on various parametric comparisons between our robot swarm and the simulation. We plan to further investigate scaling properties and are also considering a more detailed discrete event model based on a first principles approach to the robot mechatronics and control software. We believe such a model will be useful in cross comparisons between our simple emergent behaviour model(s) in studying large swarms (beyond the numbers of robots that we can afford).

Our preliminary findings suggest that for a group of 8 or more robots, the statistical or emergent behaviour is more interesting to study than that of each individual and that simulation models like this may be adequate to describe and predict collective effect such as time to complete a set goal.

It is clear that detailed models like the one we have described are wrong at a microscopic or detailed level. It is possible for our model to use very fine grained numerical integration routines and to be nearly ergodic if not precisely so. Our robot swarm is not energy conserving however and we can only reasonably expect the grosser large scale behaviour to be correct. Although the system sizes we are using (ie 8 or 16 robots) are hardly of thermodynamic scales we hope to be able to apply canonical ensemble ideas as well as micro-canonical[1].

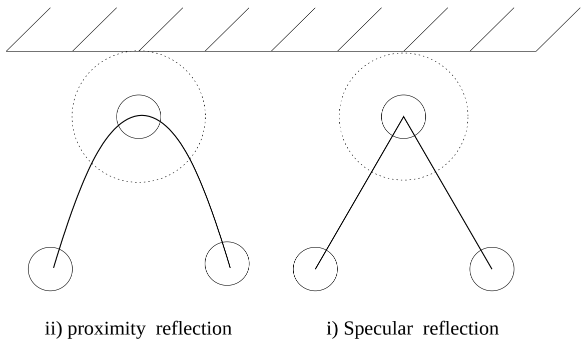

Figure 11. Robot Collisions from a wall. i) perfect specular reflection of a circular robot. ii) more realistic collision path.

In setting up the simulation model it is unclear how precise the details of the interaction between robots have to be. Figure 11 shows the obvious comparison between hard colliding disks and the locality information picked up from the robots on-board sensors. The collision/interactions of the real robots are much "softer", based on the analogue electronics used to receive information from the ultrasonics and the motor control delay/response times. In spite of this obvious flaw, the model appears capable of adequately describing clumping behaviour, and we suspect from preliminary experiments that it will usefully describe segregation behaviour in medium sized systems too.

We are also using our enhanced cybot system to investigate analogies between broadcast and multicast messaging amongst robots. We envisage a model where individual robots "change sides" dynamically during an engagement and use the wireless to proclaim their group membership rather than using fixed sensors.

We have started with simple seeking/avoidance behaviours but hope to adopt a more *ad hoc* set of heuristic models to describe cooperative behaviours such as a cooperative group seek for some resource.

Acknowledgements

The Distributed and High Performance Computing Group is a collaborative venture between the University of Wales, Bangor and the University of Adelaide in South Australia. Thanks to J.E.Story and R.G.Shepherd for assistance in setting up the Cybot robots.

References

[1] M.P. Allen and D.J. Tildesley, "Computer simulation of liquids", Clarendon Press, 1987.

[2] Bluetooth. "The Bluetooth Core Specification". Available from http://www.bluetooth.com/dev/-specifications.asp Last visited August 2002.

[3] Don Loomis, "The TINI Specification and Developers Guide", Addison Wesley Publishing co., 2001.

[4] K.A.Hawick, "Domain Growth in Alloys", PhD Thesis, Physics Department, University of Edinburgh, 1991.

[5] K.A.Hawick, H.A.James, R.G.Shepherd and J.E.Story, "Distributed Computing for Robotics", University of Wales, Bangor, Technical Note DHPC-114 www.dhpc.informatics.bangor.ac.uk/reports/114/abs-114.html

[6] J.A. Mathew, H.A. James and K.A. Hawick, "Development Routes for Message Passing Parallelism in Java," *Proc. ACM 2000 Java Grande Conference*, San Francisco, June 2000, pp 54-61.

[7] Microchip, "PICmicro microcontroller", Available from http://www.microchip.com Last visited September 2002.

[8] Cybot robot. As described in Real Robots Magazine, Published by Eaglemoss International Ltd., See also www.realrobots.co.uk.