From robot swarms to ethical robots: the challenges of verification and validation - part 1

Swarm Engineering

Alan FT Winfield

Bristol Robotics Laboratories

http://www.brl.ac.uk

RoboCheck Winter School,
University of York
1 Dec 2015

University of the
West of England
BRISTOL

This Talk

- In three parts:
  - Short introduction to Swarm Robotics
    - potential and challenges
    - flocking
  - Case Study: Adaptive Swarm Foraging
    - the algorithm
    - mathematical modelling and optimisation
  - Case Study: Reliability and Scalability
    - emergent swarm taxis
    - a reliability model

Swarm Intelligence…

– *"Any attempt to design algorithms or distributed problem-solving devices inspired by the collective behaviour of social insect colonies and other animal societies"* Bonabeau, Dorigo and Theraulaz, 1999

Termite mound

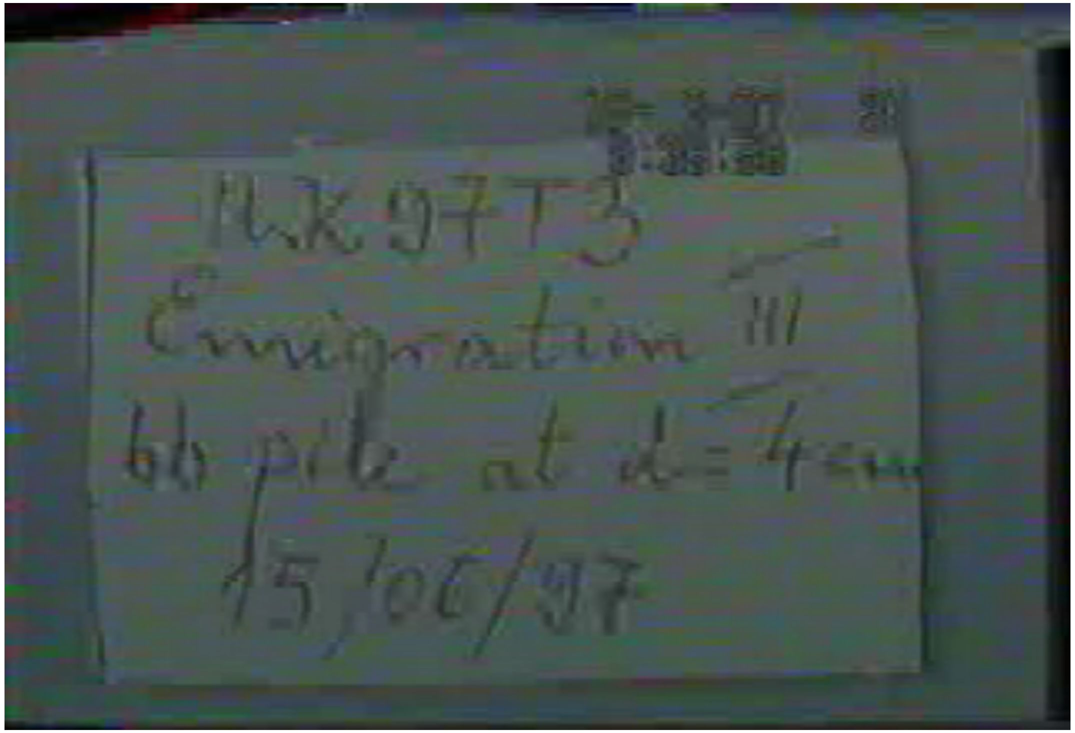

Leptothorax at work

The Potential: Swarm Robotics is characterised by...

- Relatively simple, autonomous robots
- Fully distributed, de-centralised control
  - Exploitation of agent-agent and agent-environment interaction
  - Exploitation of explicit or implicit (stigmergic) communication
  - Self-organisation and emergence
- Scalability
- Robustness

But... can we engineer solutions with swarm intelligence..?

- What are the design principles involved?
  - how do we determine the *local rules* for each individual agent, in a principled way?
- How can we validate overall behaviours that are *emergent* properties?
  - notwithstanding these (difficult) questions...
- A powerful new engineering paradigm for large scale distributed systems..?

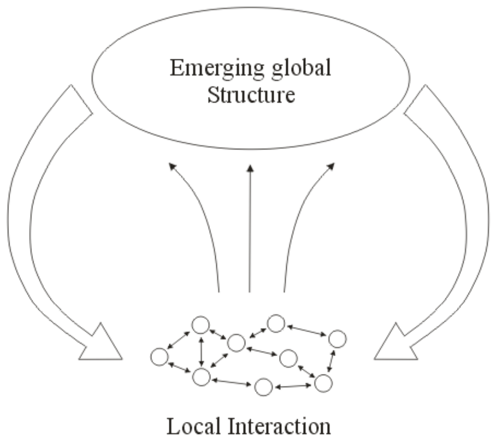

From Lewton: Complexity - Life at the Edge of Chaos

Designing the local rules

swarm = superorganism

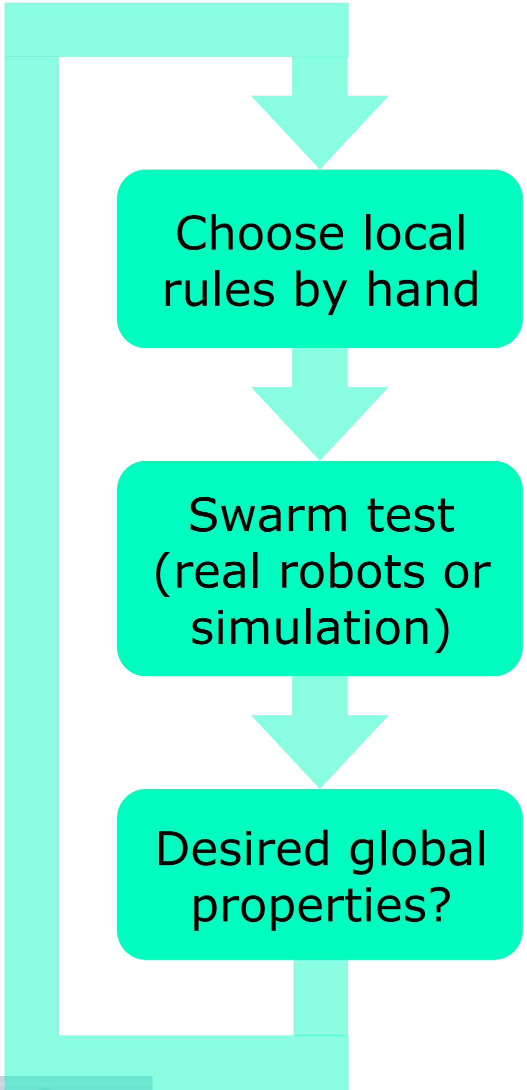

Ad-hoc
vs.
Principled approach

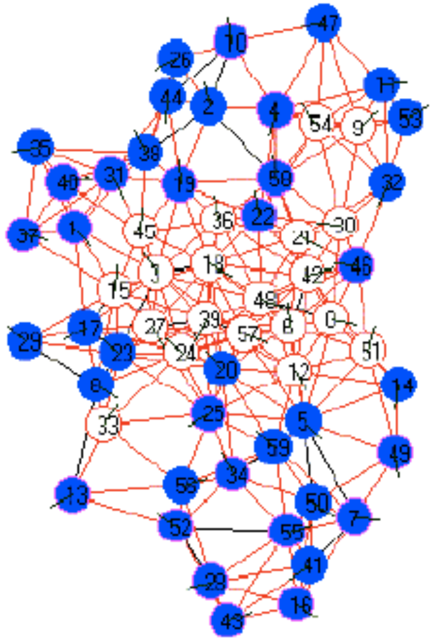

swarm = phenotype
global properties = fitness function
genotype determines local rules
Evolutionary swarm robotics

The Real-world Potential

- Any application requiring multiple distributed autonomous robots...
    - unmanned exploration/mapping/ surveying/environmental monitoring
    - robot assisted search and rescue
    - robot assisted harvesting/horticulture
    - waste processing/recycling
    - domestic or industrial cleaning
    - art and entertainment

Real-world Applications

- At the time of writing there is only one known real-world application of swarm robotics
  - A swarm of autonomous parachutes for delivering supplies
    - the *Onyx* parachutes swarm to maintain proximity so that they will not be widely dispersed on landing
    - see http://www.gizmag.com/go/6285/

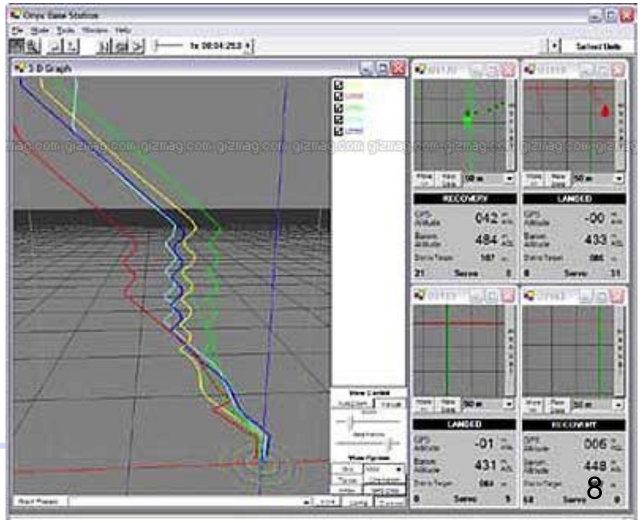

Example: the Flying Flock Project - emergent control of groups of miniature helium-filled blimps (aerobots)

A flock of Starlings

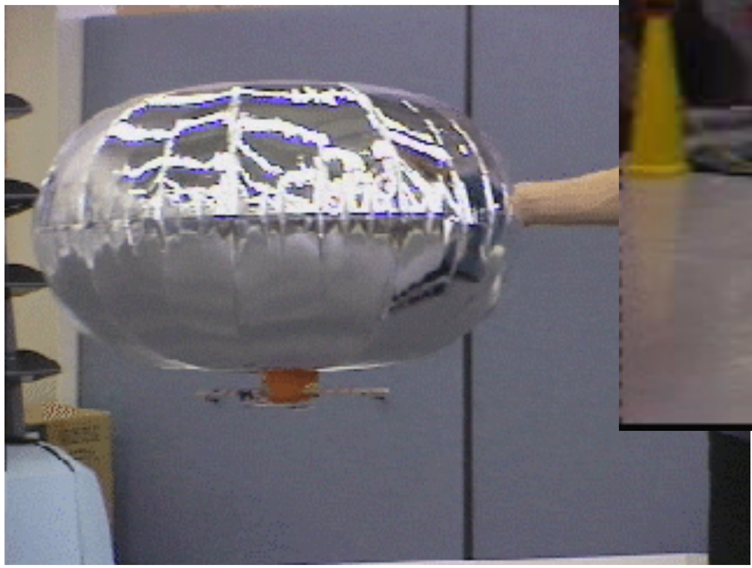

The world’s first flock of real (aero)bots in 3D [Welsby]

Case study: Foraging robots

Roomba, iRobot

Slugbot (BRL)

Zoë, Wettergreen et al, 2005

Demeter, Pilarski et al, 1999

Multi-Robot Foraging

Soda can collecting

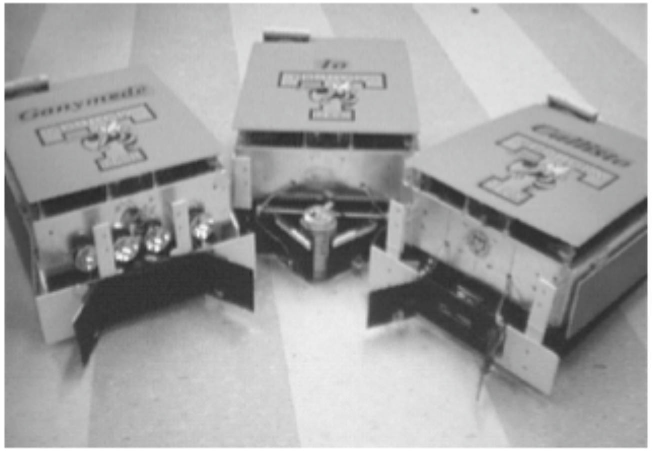

Balch et al. Io, Ganymede and Callisto: A multiagent robot trash-collecting team. AI Magazine, 16(2):39–53, 1995.

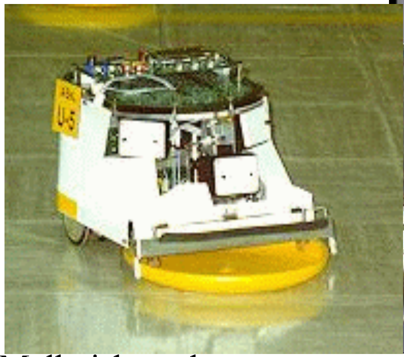

Melhuish et al.

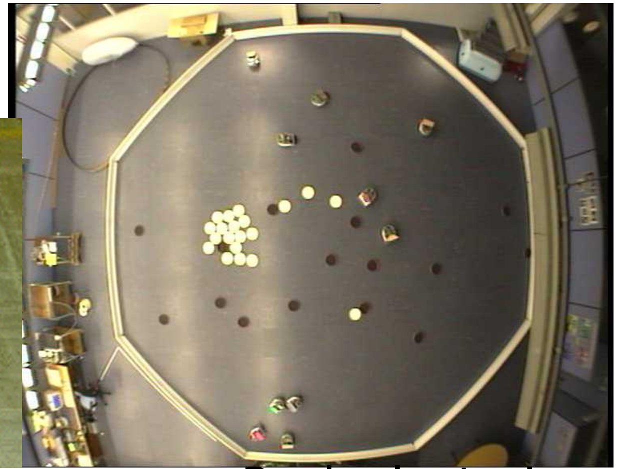

Puck clustering

Multi-robot foraging

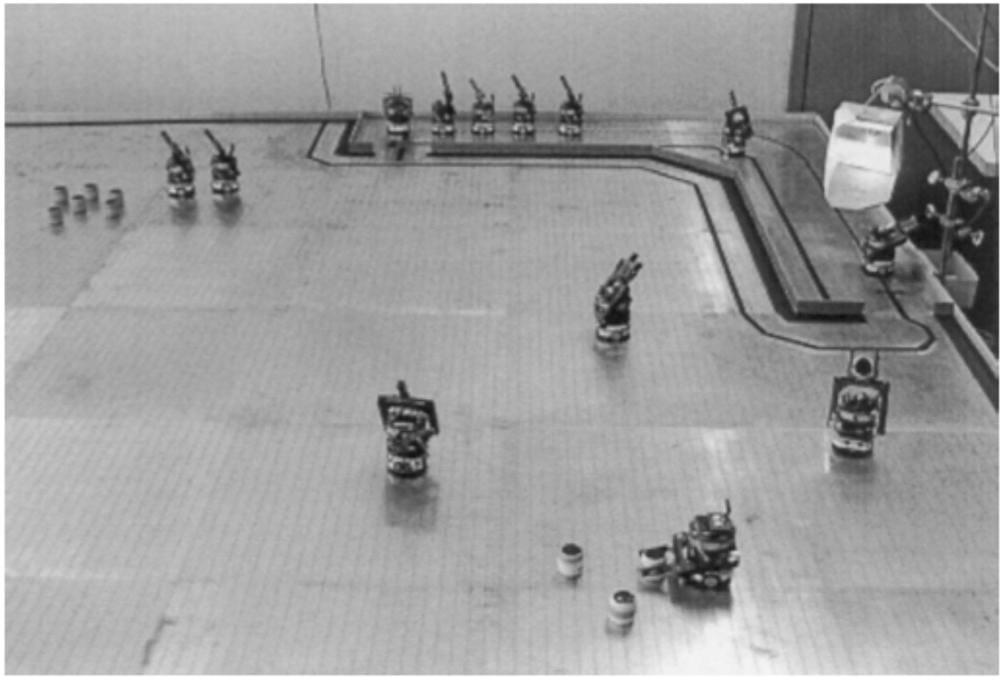

M. Krieger and J.-B. Billeter. The call of duty: Self-organised task allocation in a population of up to twelve mobile robots. Jour. of Robotics & Autonomous Systems, 30:65–84, 2000.

Multi-Robot Foraging 2

Collective manipulation

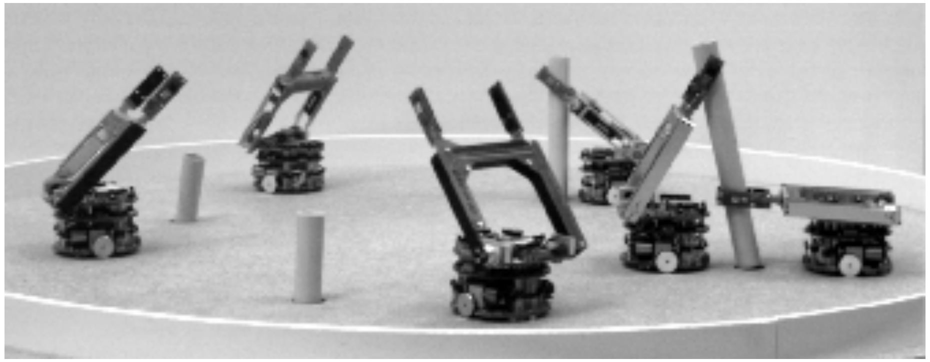

A. J. Ijspeert, A. Martinoli, A. Billard, and L. M. Gambardella. Collaboration through the exploitation of local interactions in autonomous collective robotics: The stick pulling experiment. Autonomous Robots, 11(2):149–171, 2001.

Search and Rescue, Prof Andreas Birk, Jacobs Uni, Bremen

Collective transport

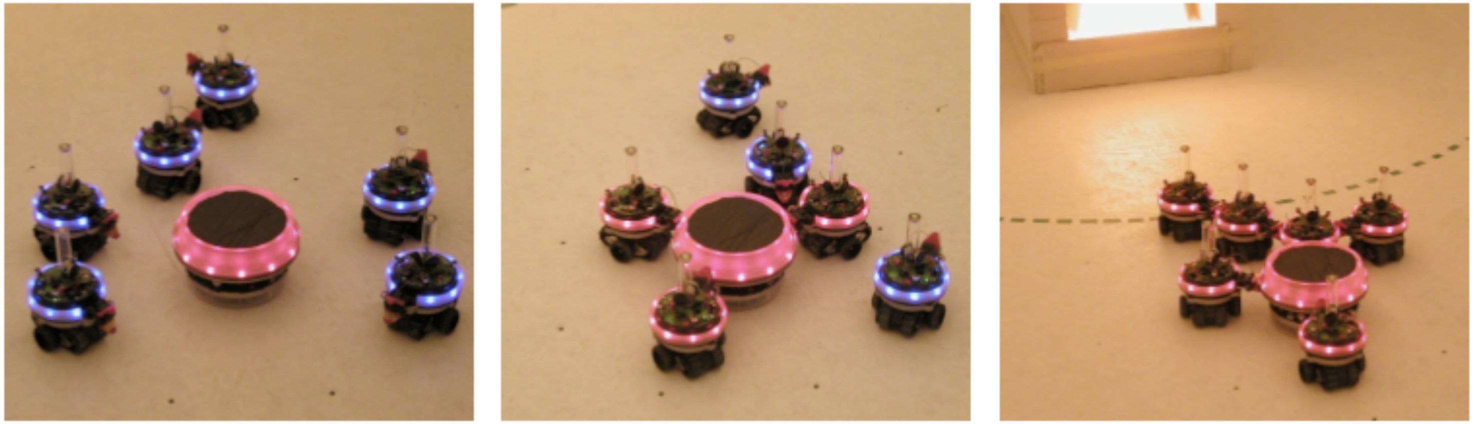

M. Dorigo, E. Tuci, T. Groß, V. Trianni, T.H. Labella, S. Nouyan, and C. Ampatzis. The SWARM-BOT pro ject. In Erol Sahin and William Spears, editors, Swarm Robotics Workshop: State-of-the-art Survey, number 3342 in Lecture Notes in Computer Science, pages 31–44, Berlin Heidelberg, 2005. Springer-Verlag

Finite State Machine for basic foraging

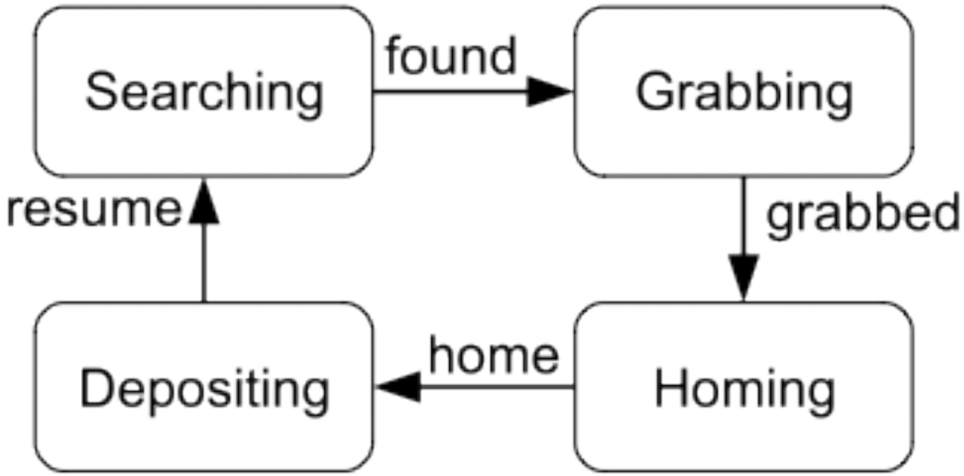

Four basic states provide an abstract model for single or multi robot foraging

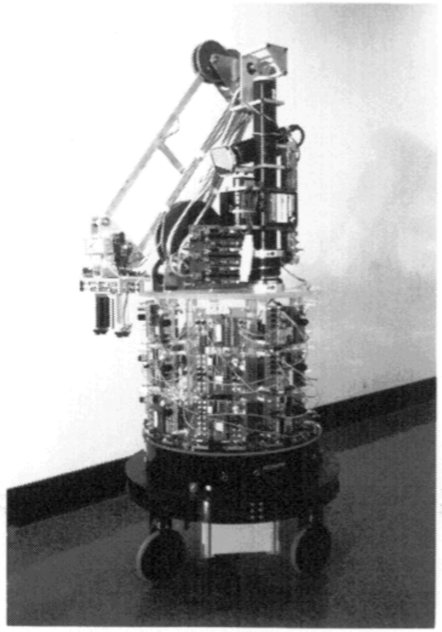

Herbert

J. H. Connell. Minimalist Mobile Robotics: A colony-style architecture for an artificial creature. Morgan Kaufmann, 1990.

Generalised FSM for foraging with division of labour

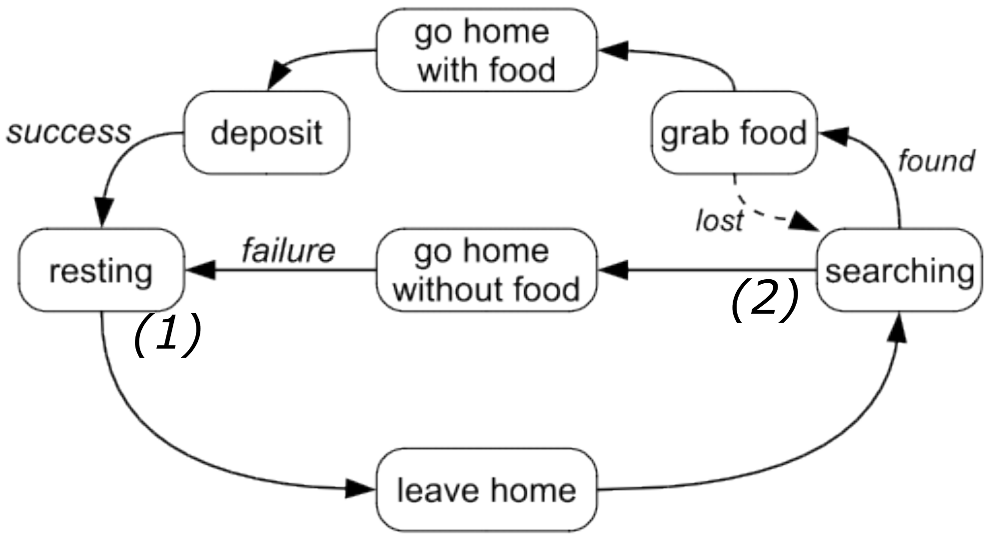

- Robots leave the nest *(1)* when some threshold condition is met
  - e.g. resting time is up or net swarm energy drops below a certain value
- Robots abandon search *(2)* when
  - e.g. searching time is up or robot energy falls below a certain value
- We seek an algorithm in which robots can locally adjust their thresholds so that the overall ratio of resters to foragers adapts to the amount of food in the environment

Note: ‘food’ is a metaphor for any objects to be collected

Energy foraging

- Consider the special case of multi-robot foraging in which robots are foraging for their own energy. For an individual robot foraging costs energy, whereas resting conserves energy.
  - Each robot consumes energy at A units per second while searching or retrieving and B units per second while resting, where $A > B$
  - Each discrete food item collected by a robot provides C units of energy to the swarm
  - The average food item retrieval time, is a function of the number of foraging robots x, and the density of food items in the environment, ρ, thus $t = f(x, \rho)$

Strategies for cooperation

- Each robot has a search time threshold T_s and a rest time threshold T_r
  - Internal cues. If a robot successfully finds food it will reduce its T_r; conversely if the robot fails to find food it will increase its T_r
  - Environment cues. If a robot collides with another robot while searching, it will reduce its T_s and increase its Tr times
  - Social cues. When a robot returns to the nest it will communicate its food retrieval success or failure to the other robots in the nest. A successful retrieval will cause the other robots in the nest to increase their T_s and reduce their T_r times. Conversely failure will cause the other robots in the nest to reduce their T_s and increase their T_r times

	internal cues	social cues	environment cues
S_1 (baseline)	×	×	×
S_2	✓	×	×
S_3	✓	✓	×
S_4	✓	✓	✓

Adaptive foraging with changing food density

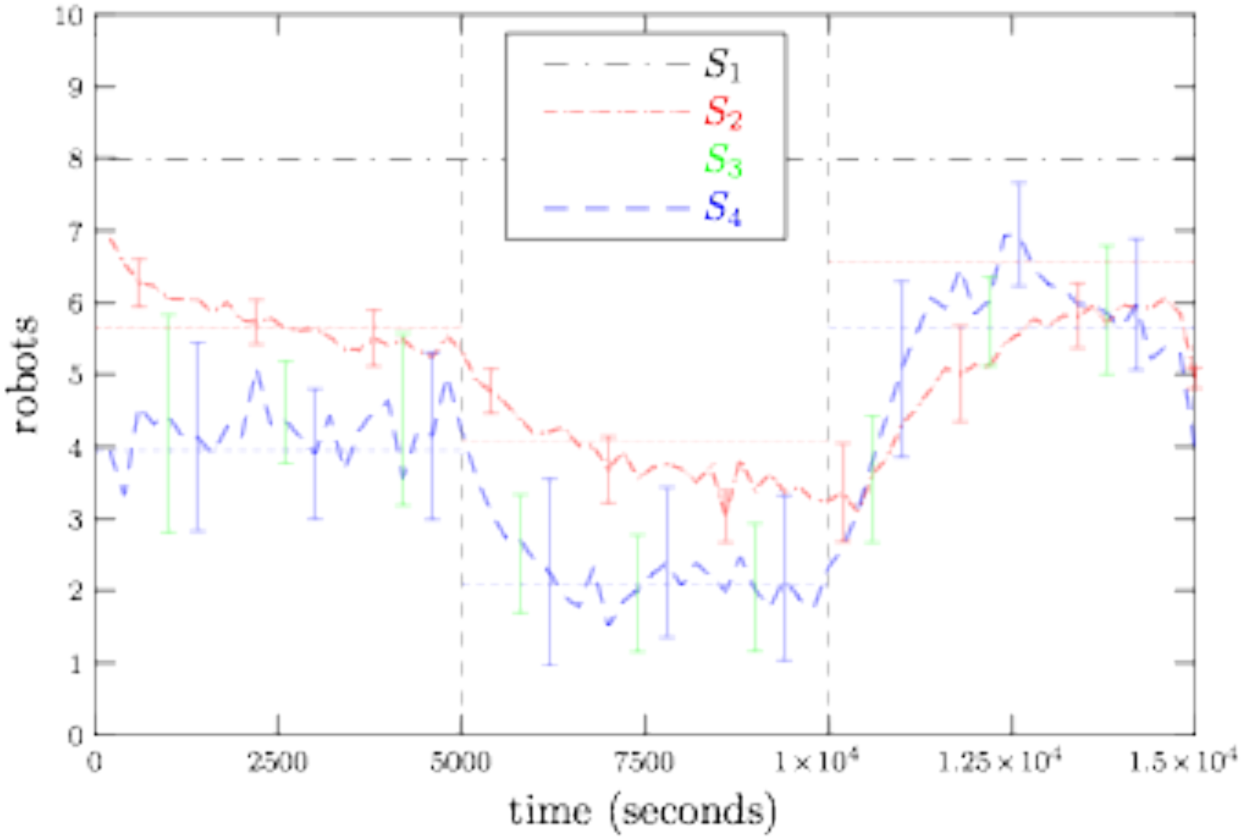

W. Liu, A. F. T. Winfield, J. Sa, J. Chen, and L. Dou. Towards energy optimisation: Emergent task allocation in a swarm of foraging robots. Adaptive Behaviour, 15(3):289– 305, 2007.

Number of foraging robots *x* in a foraging swarm of *N* = 8 robots. S1 is the baseline (no cooperation strategy); S2, S3 and S4 are the three different cooperation strategies. Food density changes from 0.03 (medium) to 0.015 (poor) at *t* = 5000, then from 0.015 (poor) to 0.045 (rich) at *t* = 10000. Each plot is the average of 10 runs.

Mathematical Modelling

- We model apply the probabilistic approach of Martinoli *et al**.
- We take the Finite State Machine (FSM)
  – express as an ensemble of probabilistic FSMs...which lead to a set of difference equations
  – geometrically estimate the transition probabilities
  – compare the model with experimental data

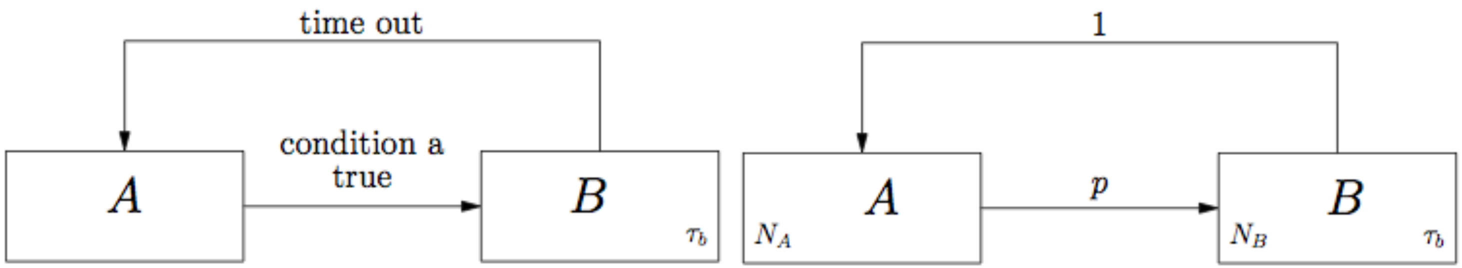

(a) finite state machine (b) PFSM

*See e.g. Martinoli, Easton and Agassounon, IJRR 23(4), 2004

Developing a mathematical model

Finite State Machine ⇨ Probabilistic Finite State Machine (PFSM)*

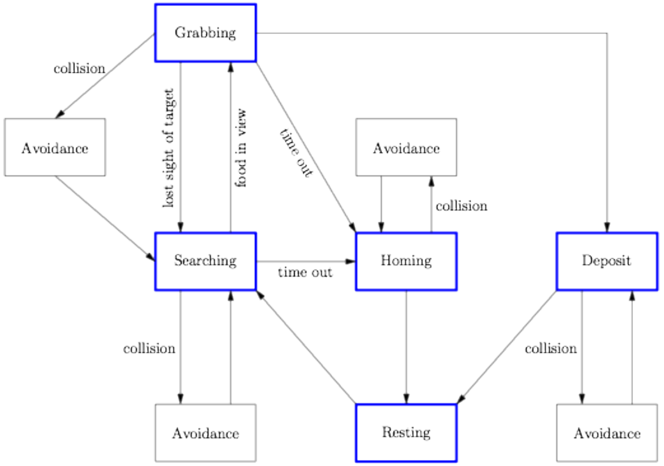

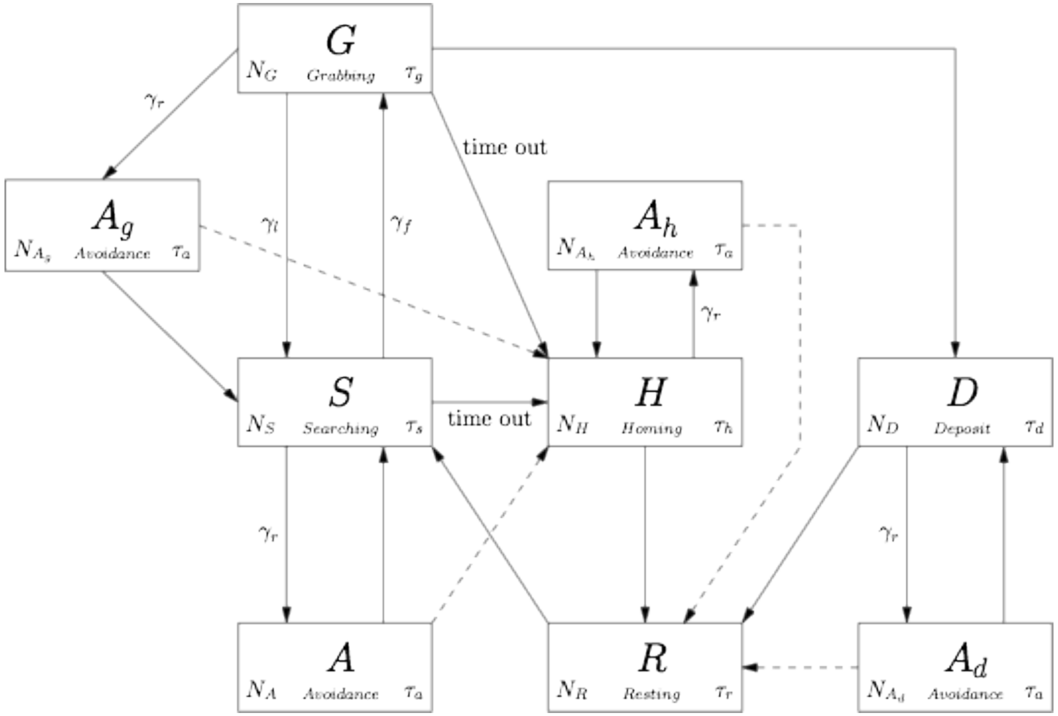

PFSM parameters:

$N_.$ number of robots in state .

$\tau_.$ time in state .

γ_f probability of finding food

γ_l probability of losing it

γ_r probability of collision

Difference equations

- For the PFSM we next develop a set of difference equations, e.g.

$$
\begin{aligned}
N_S(k+1) = & N_S(k) + \gamma_l(k) N_G(k) + \Delta_R(k - T_r) + \left[\Delta_A(k - T_a) - \Omega_A(k - T_a)\right] \\
& + \left[\Delta_{A_g}(k - T_a) - \Omega_{A_g}(k - T_a)\right] - \gamma_r(k) N_S(k) - \gamma_f M(k) N_S(k) \\
& - \Gamma_S(k+1)
\end{aligned}
$$

This appears complex because of multiple sampling rates and different priorities of behaviours

Geometrical estimation of state transition probabilities

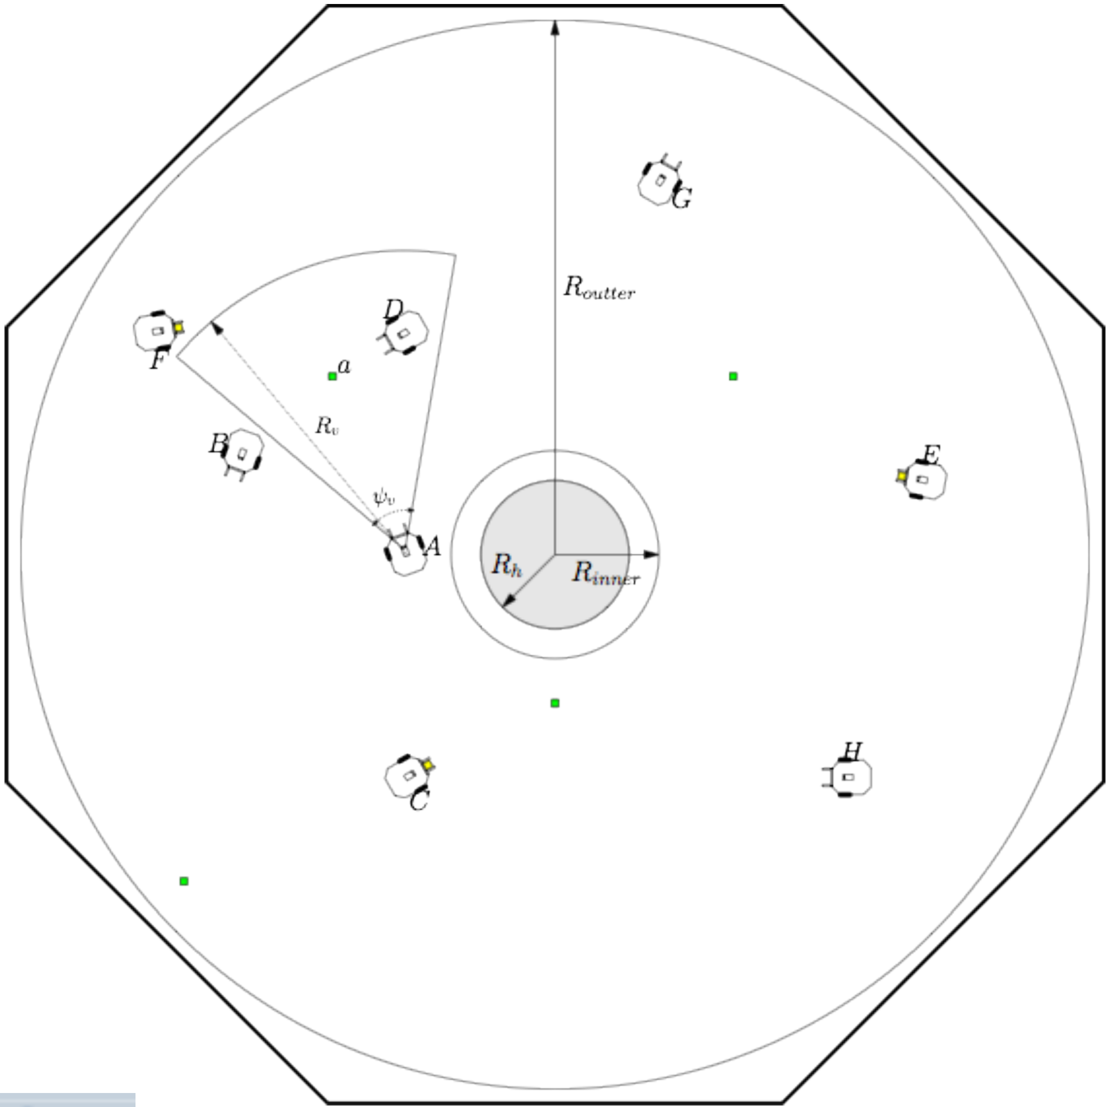

- Three simplifying assumptions:
  - place a circular nest at the centre of a circular arena
  - food items are uniformly distributed
  - robots have an equal probability of occupying any position in the arena
  - the relative heading between any two robots varies uniformly in the range 0° to 360°

probability of finding a food item: γ_f

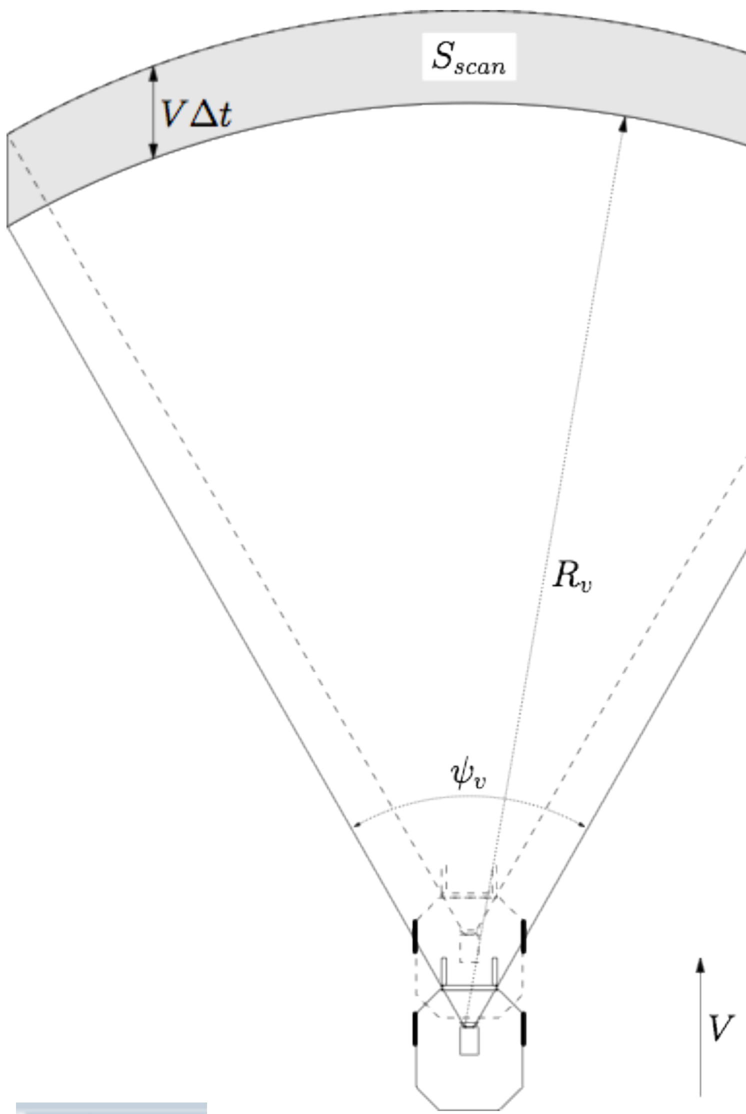

Probability to find 1 food item:

$$P_f = \frac{S_{scan}}{S_f} = \frac{\psi_v R_v V \Delta t}{\pi(R_{outer}^2 - R_{inner}^2)}$$

To find at least 1 of *M(k)* food items:

$$\gamma_f(k) = 1 - (1 - P_f)^{M(k)} \approx P_f M(k)$$

probability of losing a food item: γ_l

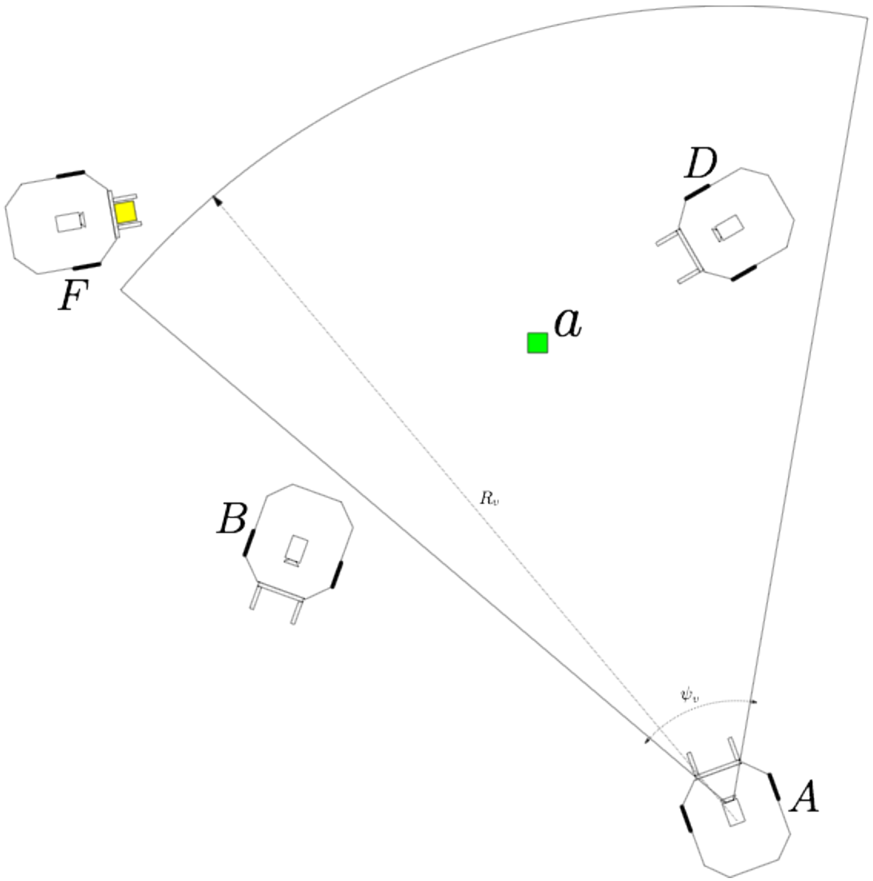

Robot *A* will lose food item *a* if:
A is not the closest to *a*, and
at least one other robot moves to *a*

Probability of losing food item *a*

$$\left(1 - \frac{1}{N_{f_a}}\right)\left(1 - (1 - \frac{p_g}{M_{fa}})^{N_{fa}-1}\right)$$

collision probability: γ_r

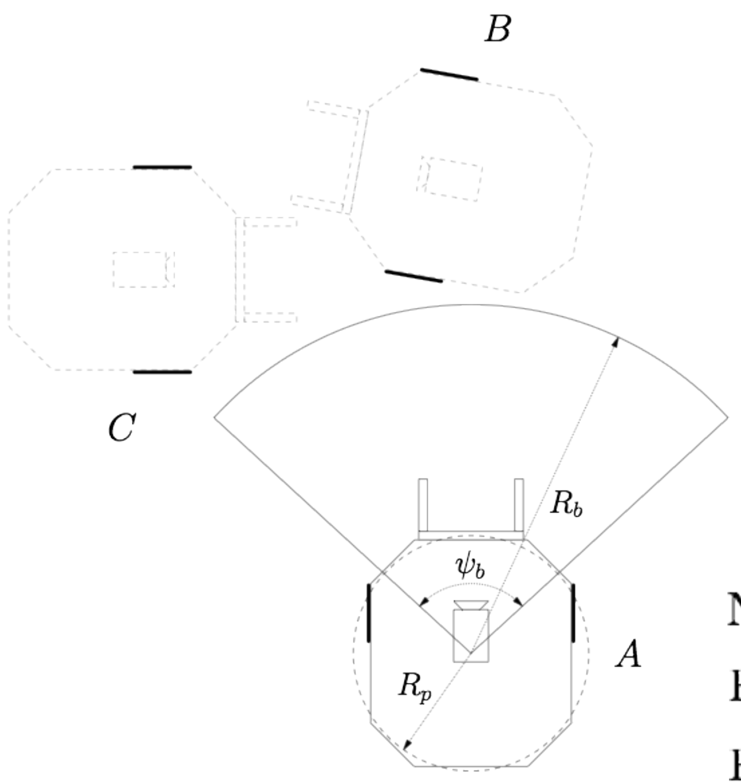

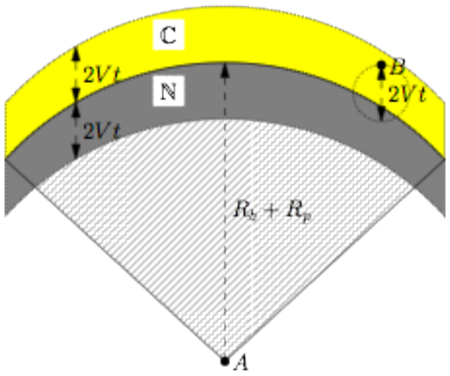

$\mathrm{N}_{active}(k)$ – number of robots outside the nest

P_{in} – the probability that a robot is in area ℂ

P_a – the probability that a robot in area ℂ will move to ℕ

$$\gamma_r(k+1) = 1 - (1 - P_{in}P_a)^{N_{active}(k)-1}$$

Estimation of time parameter τ_g

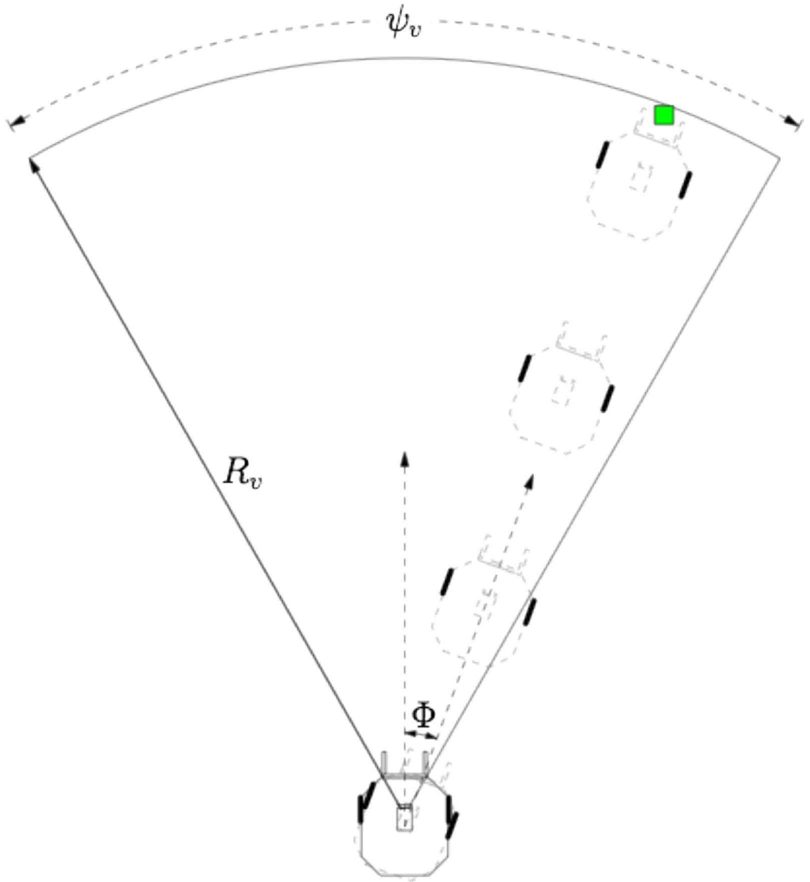

When a food item is in view the robot needs to

1. turn to face the food
2. move forward until close enough to grab it
3. grab and lift it

Average grabbing time:

$$\tau_g = \frac{\psi_v}{2w_1} + \frac{R_v}{V} + t_l$$

Validation of the model

Sensor based simulation calibrate and validated by rea robot measurement Using Player/Stage.

parameters for simulation (Player/Stage)

parameters	value	parameters	value
V	$0.15m/s$	R_{outer}	$3m$
w_1	$15°/s$	E_r	$1unit$
w_2	$15°/s$	αE_r	$10units$
ψ_v	$60°$	E_c	$2000units$
ψ_b	$95°$	Δt	$0.25sec$
R_v	$2m$	t_l	$2sec$
R_b	$0.4m$	τ_a	$2sec$
R_p	$0.13m$	τ_s	$100sec$
R_h	$0.5m$	τ_r	$[0, 200]sec$
R_{inner}	$0.7m$	p_{new}	0.04

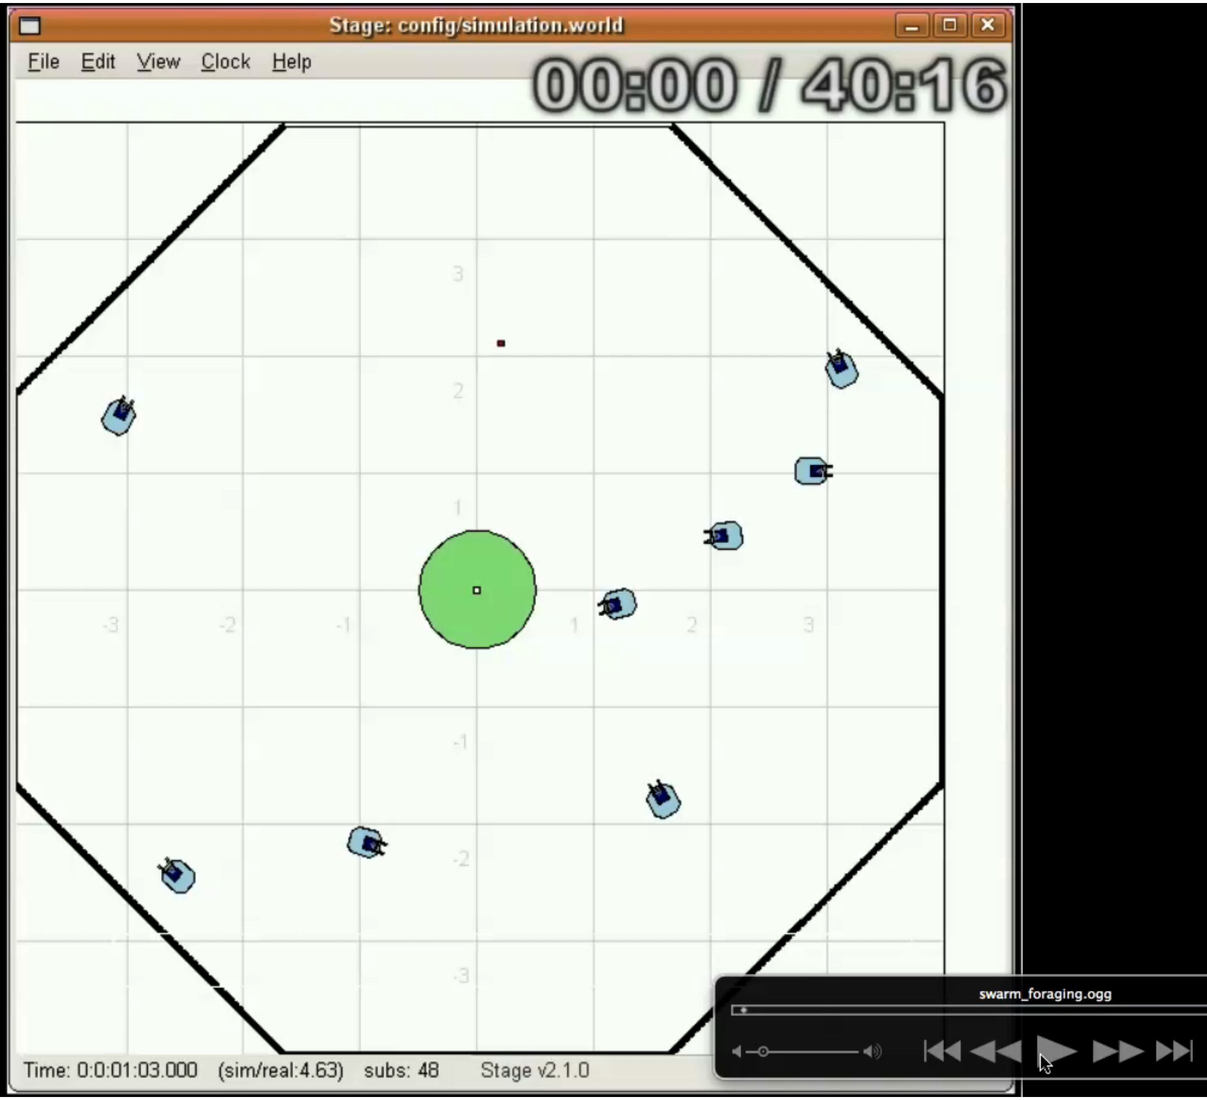

Robot platform

- Experimental platform: the LinuxBot*

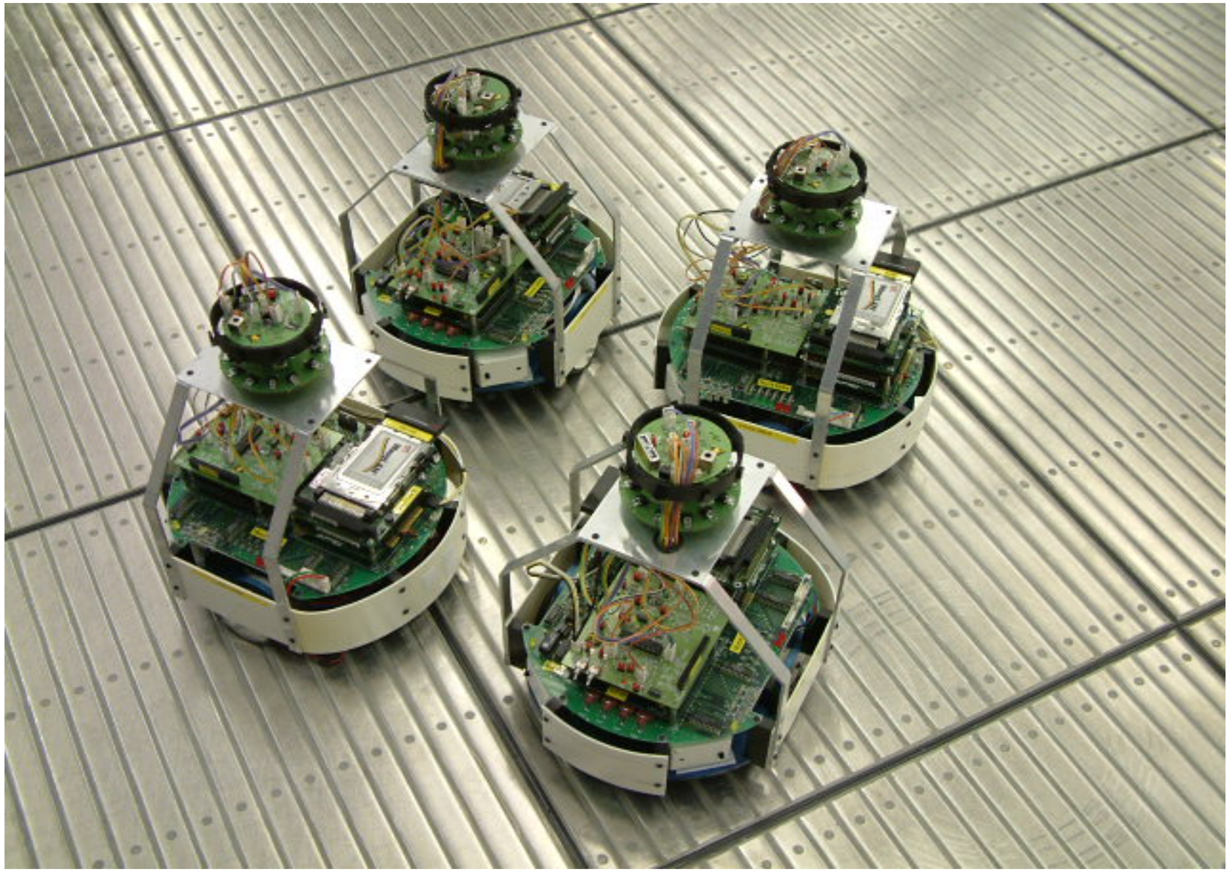

Model calibration

*See: Winfield & Holland, Microprocessors & Microsystems 23(10), 2000.

validation of the model (2)

Net swarm energy, (left) varying resting time threshold τ_r, (right) for τ_r = 80s

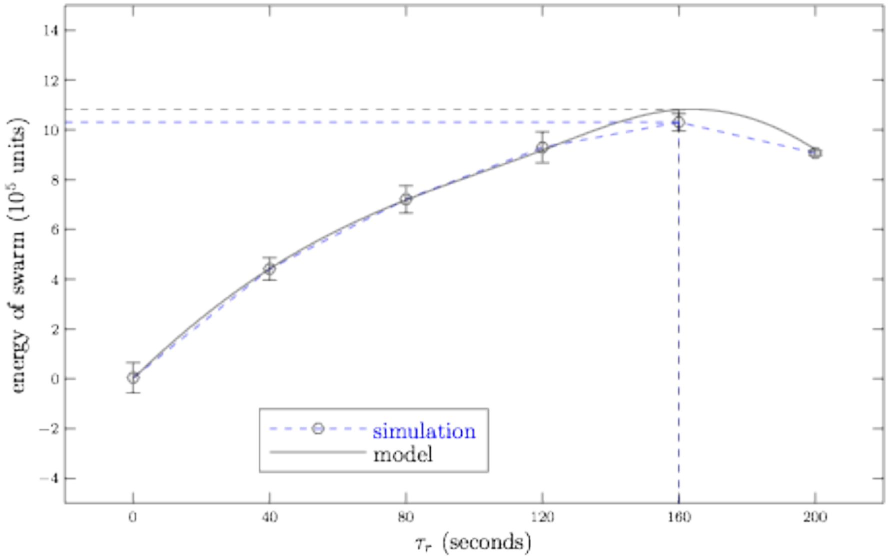

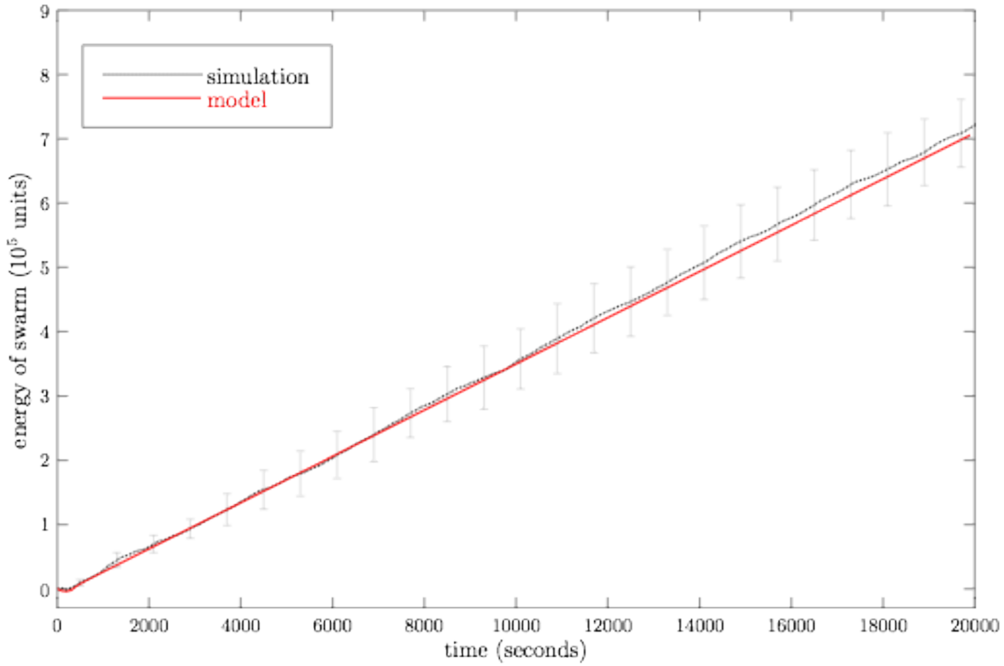

$$E(k+1) = E(k) + E_c \Delta_D (k - T_d) - E_r N_R(k) - \alpha E_r (N_0 - N_R(k))$$

validation of the model (3)

Average number of robots in states *searching*, *resting* and *homing* for τ_r= 80s

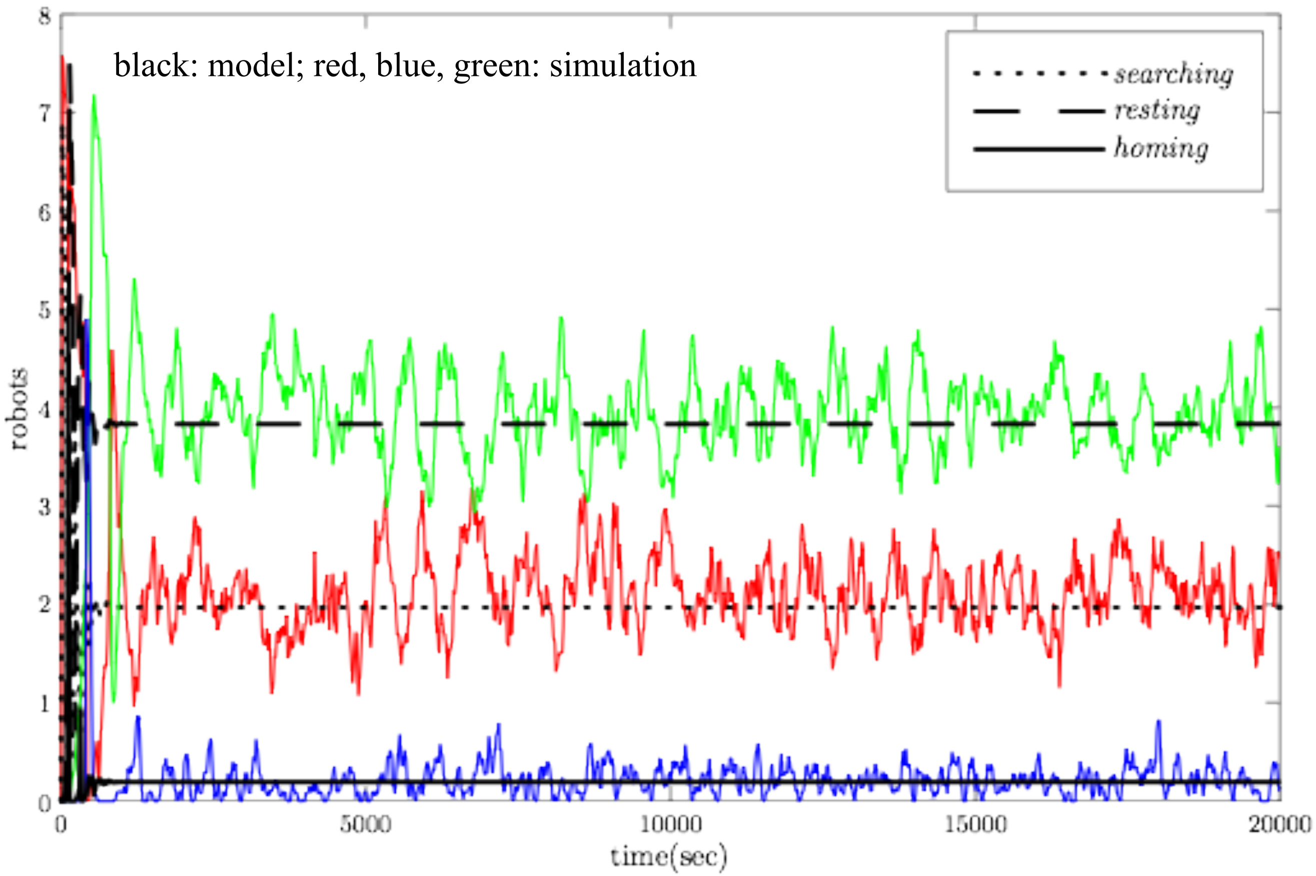

Liu W, Winfield AFT and Sa J, 'Modelling Swarm Robotic Systems: A Case Study in Collective Foraging', Proc. Towards Autonomous Robotic Systems (TAROS 2007), pp 25-32, Aberystwyth, 3-5 September 2007.

Extend the model to *adaptive* foraging

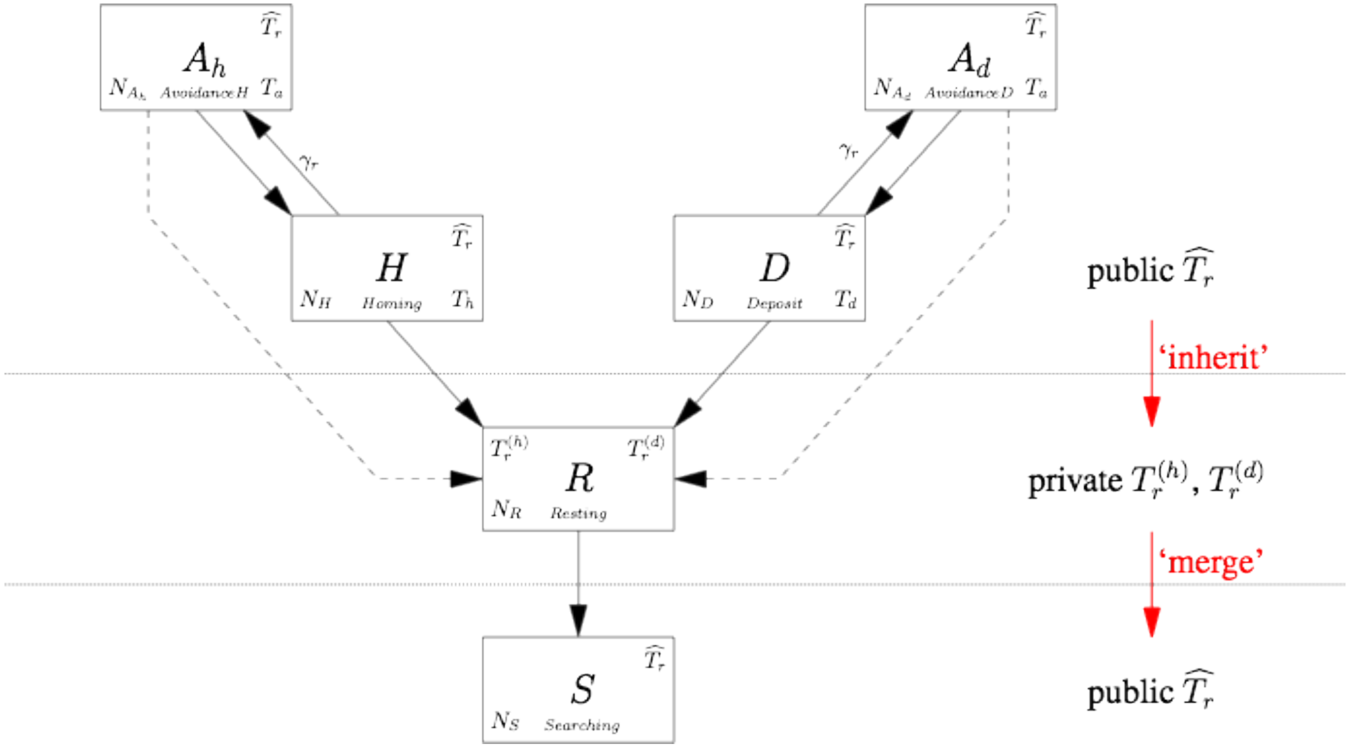

We introduce the concept of short time lived sub-PFSMs, with 'private' parameters

Model of adaptive foraging: validation of the model

Variable food density: 0.45, 0.4, 0.35

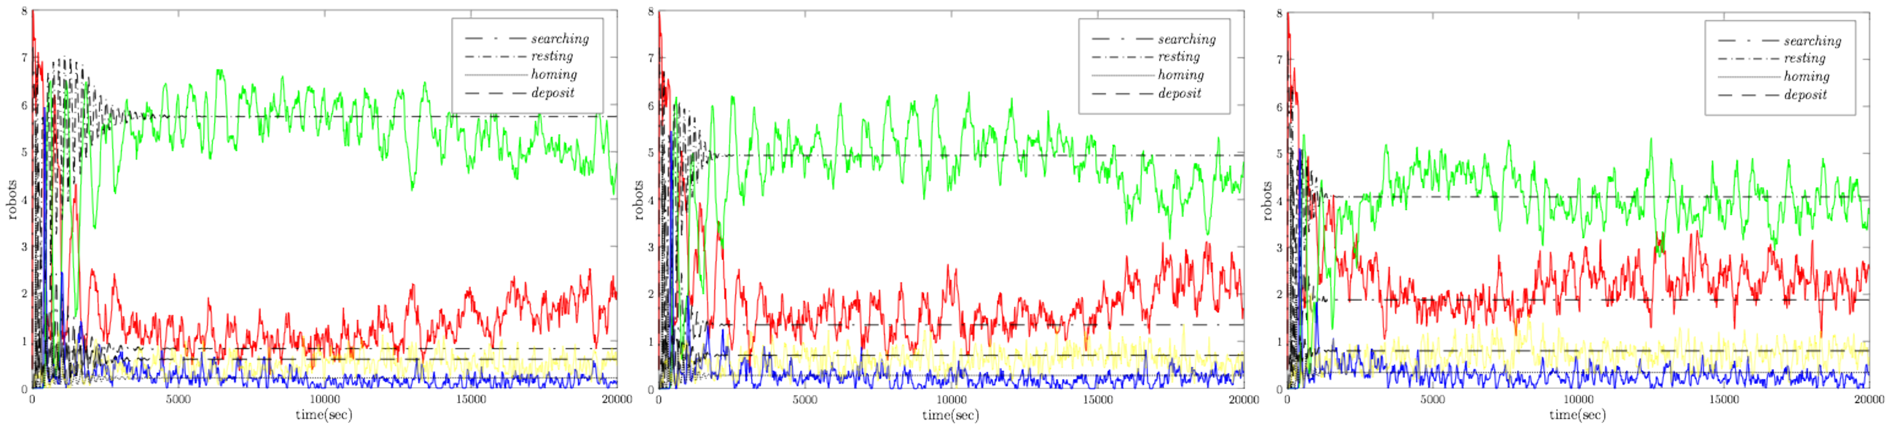

Liu W, and Winfield AFT, 'A Macroscopic Probabilistic Model for Collective Foraging with Adaptation', International Journal of Robotics Research, doi:10.1177/0278364910375139.

We were then able to use this model, together with a real-coded GA, to optimise the adjustment factors
these are the precise amounts by which the time thresholds are increased or decreased by the internal, social or environmental 'cues'

Case study: emergent swarm taxis

- A minimalist approach
- aggregation:
  - short range: obstacle avoidance (repulsion)
  - longer range: maintain number of connected neighbours (attraction)
- and beacon taxis:
  - see next slide
- Note swarm behaviour requires *team working*

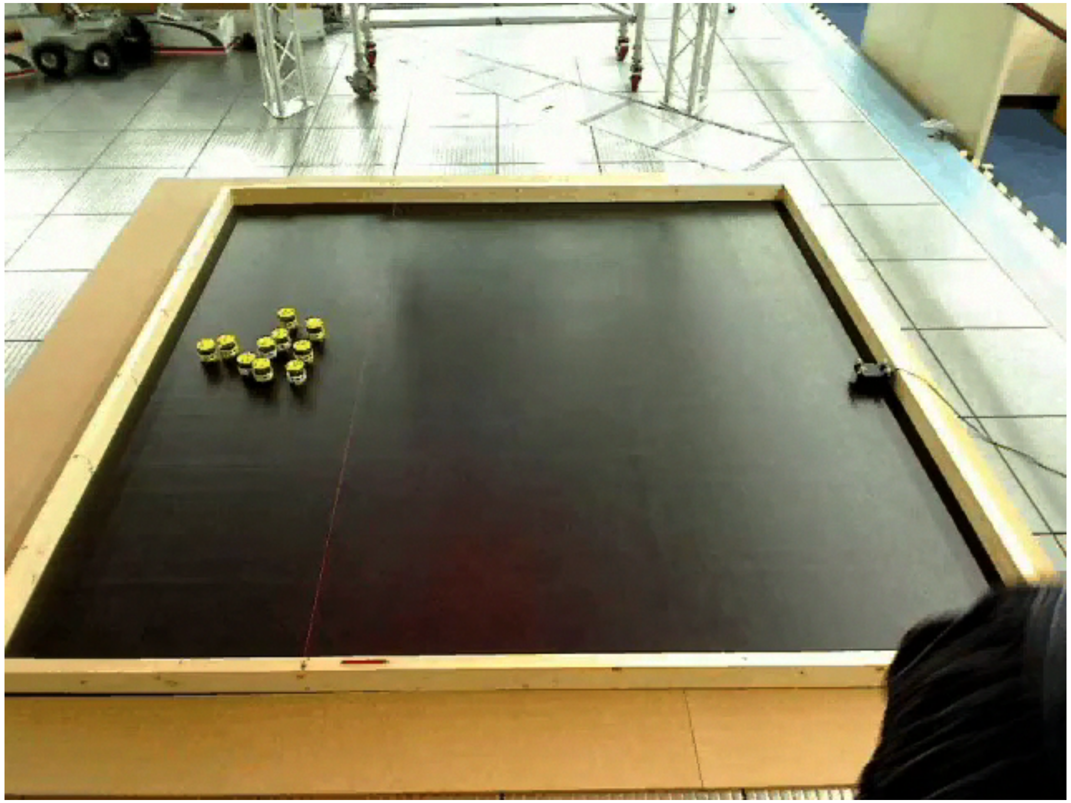

10 robots, IR beacon on the right, 25x speedup

Symmetry breaking leads to swarm taxis

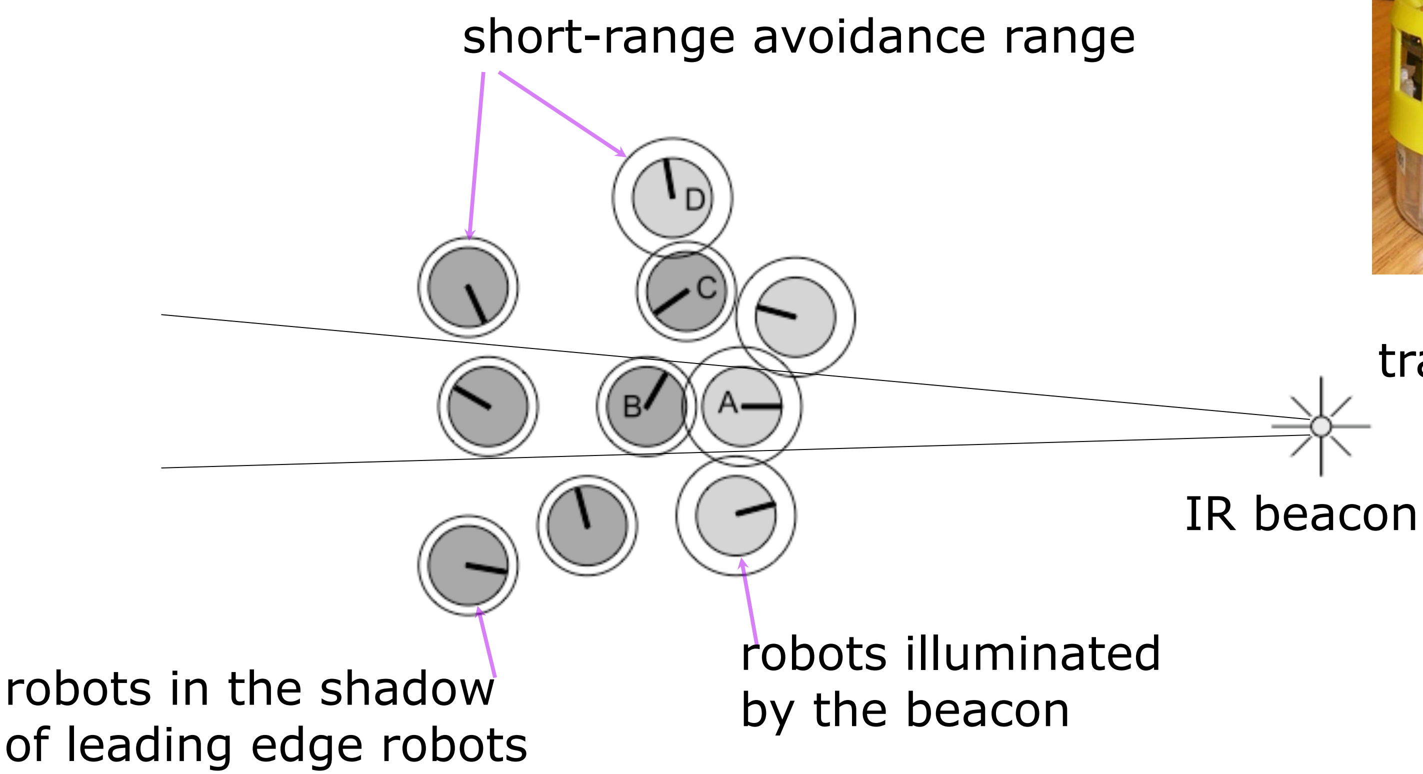

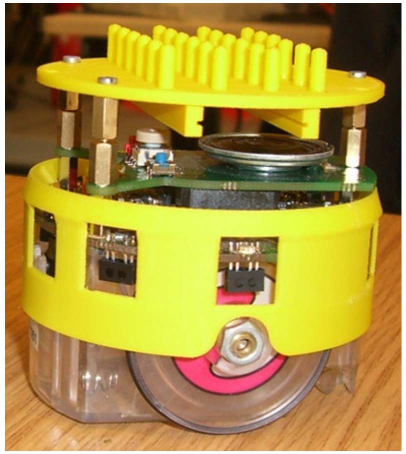

e-puck with tracking hat and skirt

The k-out-of-N:G reliability model

The probability that at least k out of N robots are working at time t:

$$P(k,N,t) = \sum_{i=k}^{N} \binom{N}{i} (e^{-t\lambda})^i (1-e^{-t\lambda})^{N-i} \qquad \lambda = \frac{1}{MTBF}$$

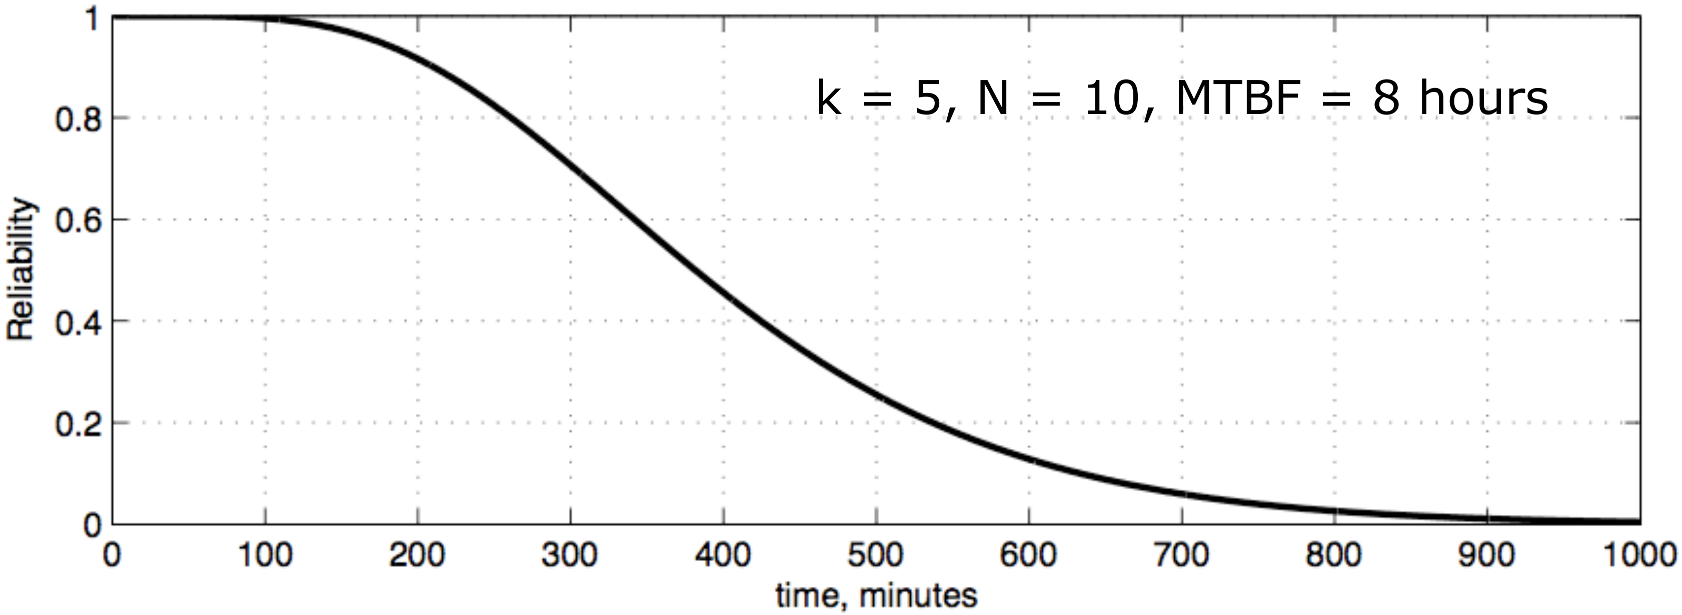

Failure modes analysis

- Case 1: *complete failures of individual robots*
  - failed robots become static obstacles in the environment
- Case 2: *failure of a robot's IR sensors*
  - failed robots leave the swarm and become dynamic obstacles in the environment
- Case 3: *failure of a robot's motors only*
  - failed robots have the effect of ***anchoring*** the swarm

Induce worst-case partially failed robots

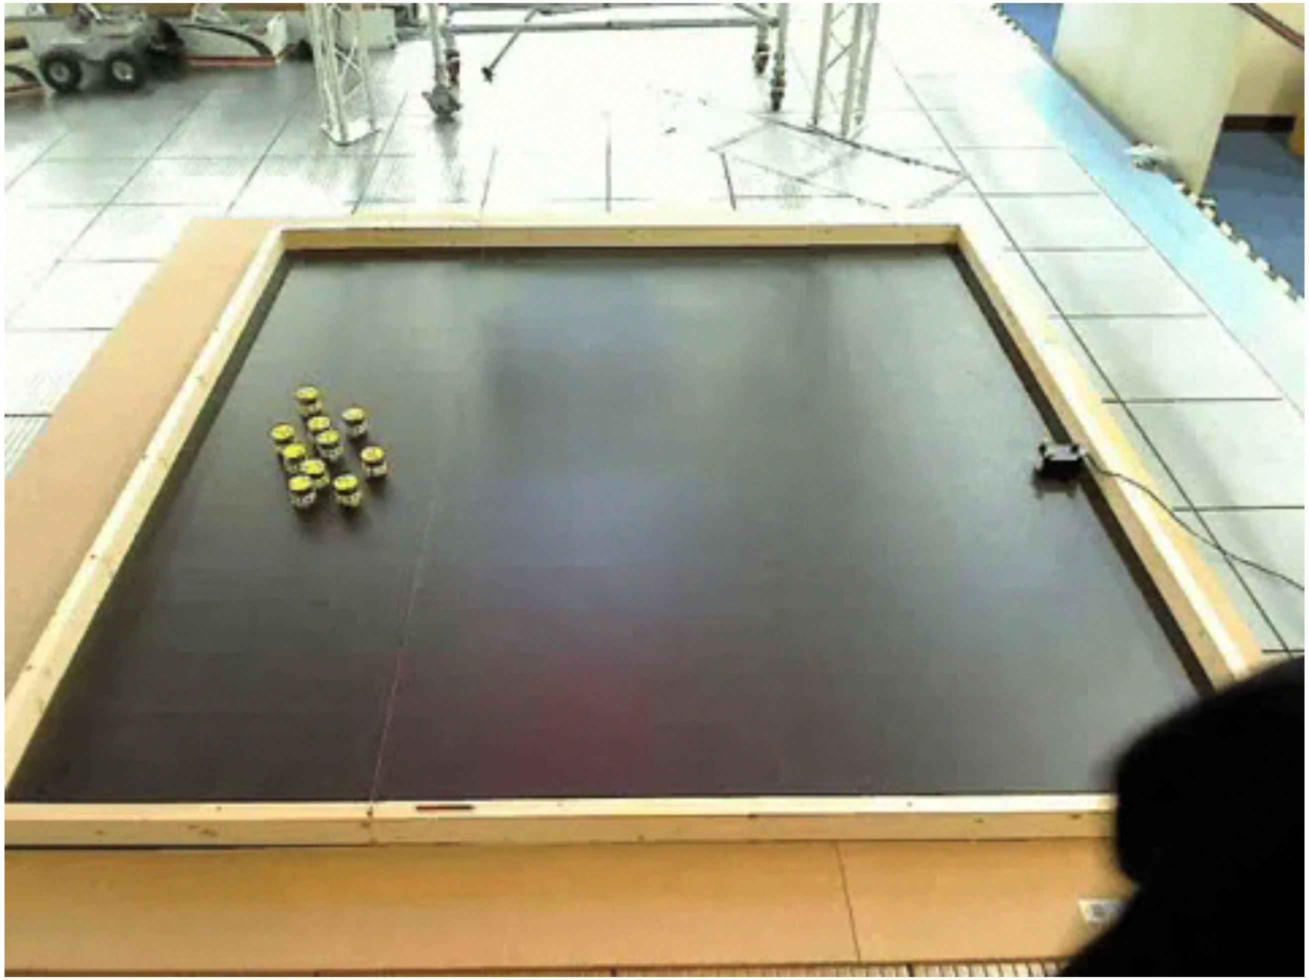

2 simultaneous case 3 partial failures

Introduce the notion of swarm self-repair

Case 1

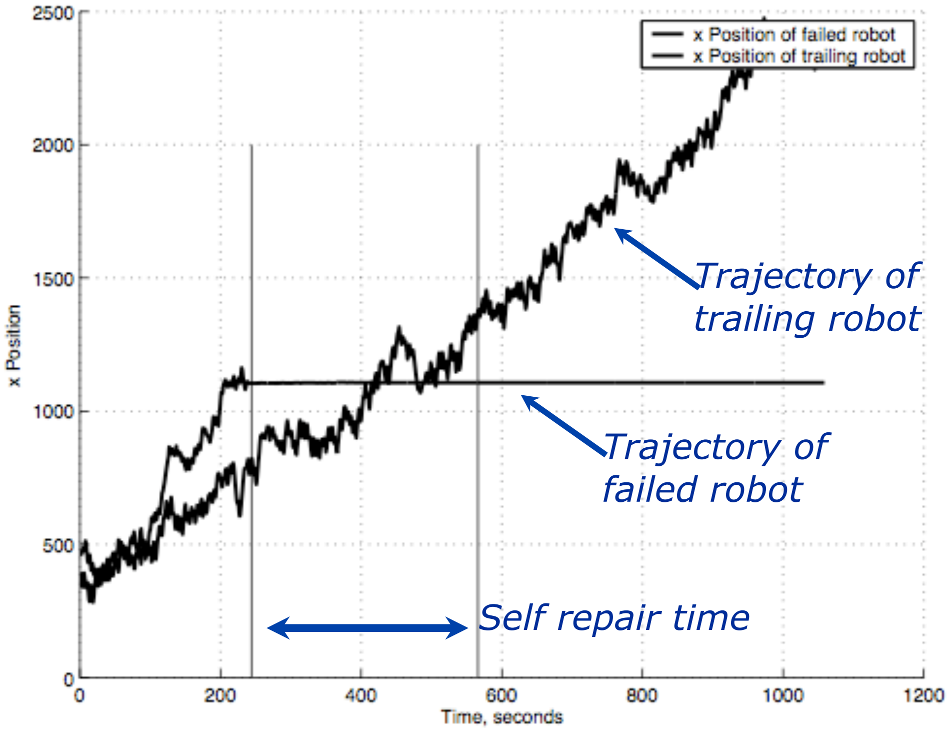

Single robot complete failure

Case 3

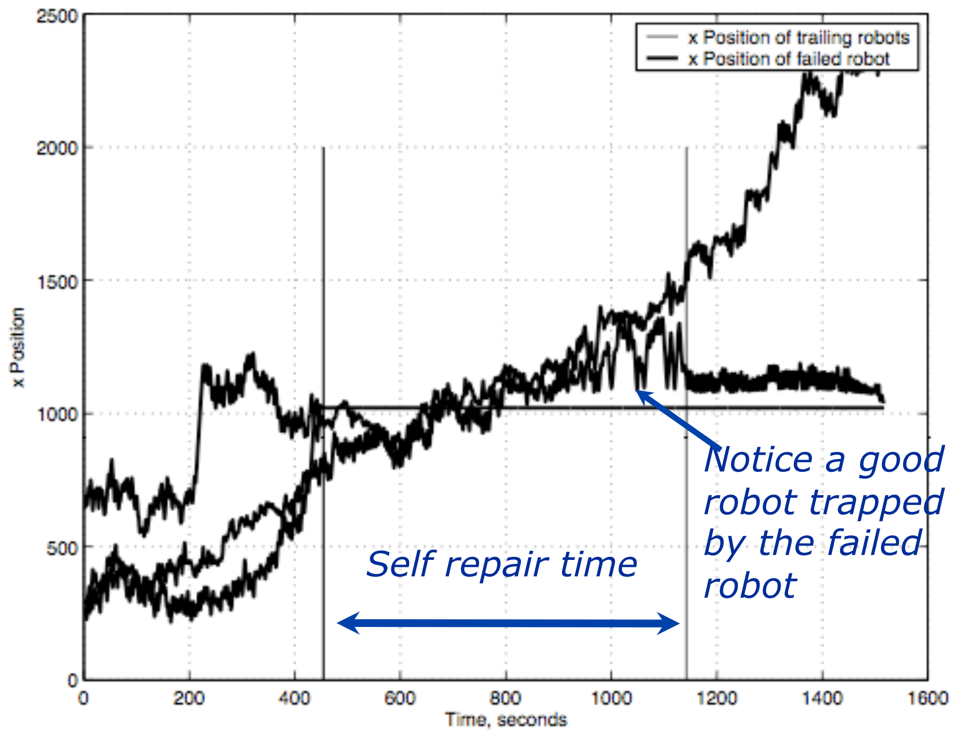

Single robot partial failure

Mean swarm self-repair times

Table 1 Mean swarm self-repair times for the case study swarm of $N = 10$ e-puck robots. Ten runs for each case. *Here the swarm reached the beacon in only 6 of 10 runs.

Case	Mean (s)	Std. Dev. (s)
Baseline (no penalty)	328	174
One failed robot Case 1	387	132
Two failed robots Case 1	453	172
One failed robot Case 3	879	417
Two failed robots Case 3	1279	see note*

Estimate k for case 3 partial failure

- Conservatively $k = 0.9N$
  - in other words, we believe the swarm can tolerate 10% of case 3 failures at any one time (i.e. within the swarm self-repair time)

Estimate swarm self-repair time

Since a robot can fail anywhere in the swarm the average distance the swarm needs to move to escape the failed robot is half the diameter of the swarm, i.e. *t = d/2v, d = swarm diameter, v = swarm velocity*

We know

$$v(N) = CN^{-\frac{1}{2}} \quad \text{and} \quad d(N) = D\sqrt{N}$$

Thus

$$t(N) = \frac{D}{2C}N$$

Therefore swarm self repair time t is linear with N.

With *N=10* and 1 partially failed robot mean swarm self repair time was measure as *870s*, thus the constant S = *D/2C = 87.9*

Reliability as a function of swarm size for swarm with partial failures

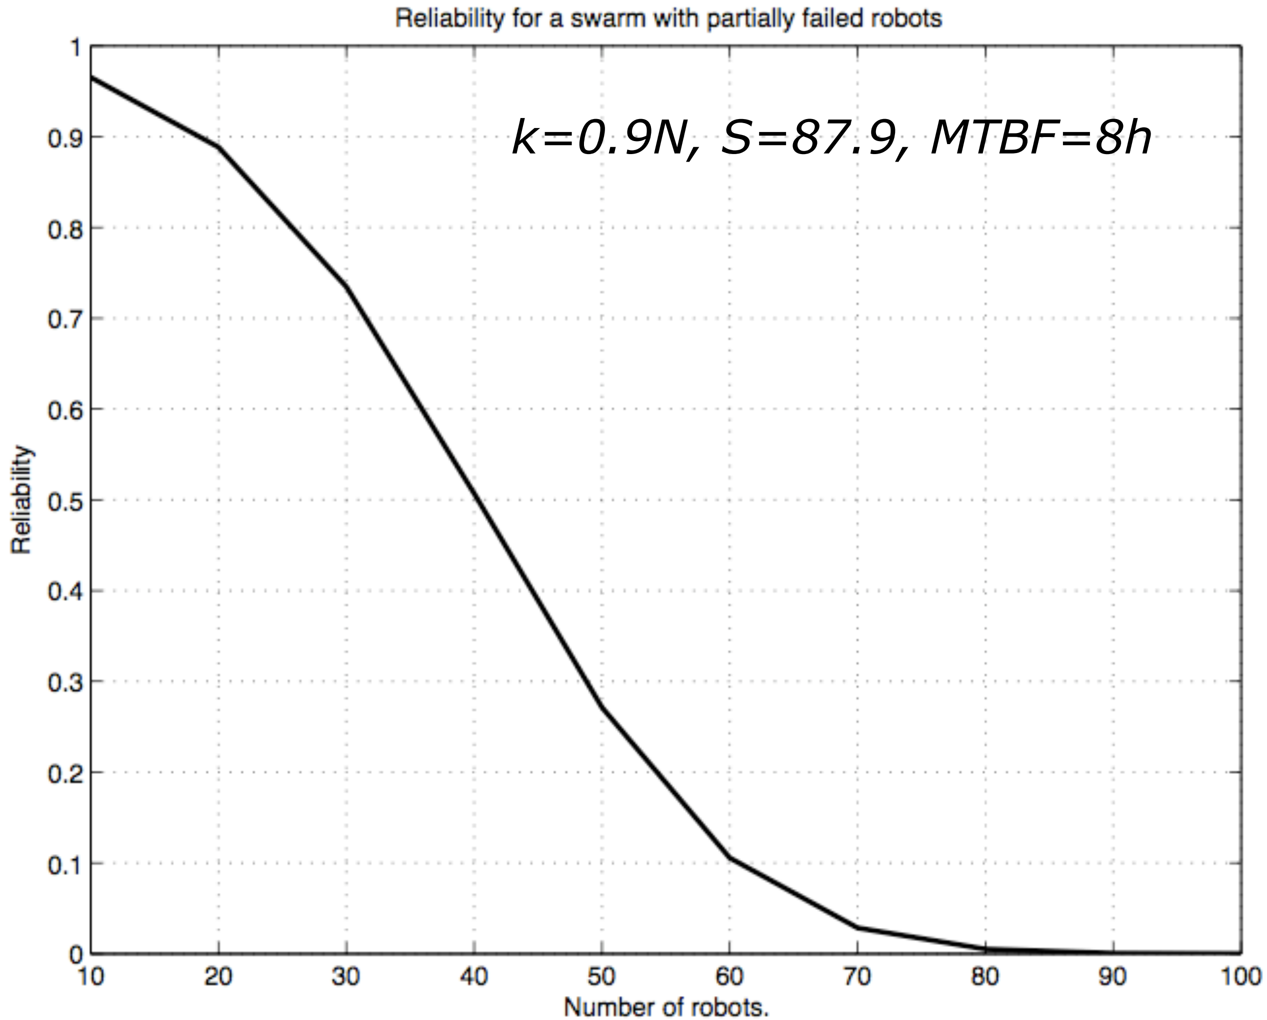

Discussion

- We need to revise our assumptions of swarm robustness and scalability
  - but note that swarms do still have a high degree of fault tolerance
- This work strongly suggests that large-scale swarms (which rely on emergence or self-organising mechanisms) will require more sophisticated active internal mechanisms for dealing with worst-case partial failures:
  - i.e. an *immune system*

See: Bjerknes JD and Winfield AFT, 'On Fault-tolerance and Scalability of Swarm Robotic Systems', in Proc. Distributed Autonomous Robotic Systems (DARS 2010)

Thank you!

EPSRC
Engineering and Physical Sciences Research Council

SEVENTH FRAMEWORK PROGRAMME

- Acknowledgements, colleagues in the BRL, but especially:
  - Dr Chris Harper, Dr Julien Nembrini, Dr Wenguo Liu, Dr Jan Dyre Bjerknes

- Further relevant publications:

  - AFT Winfield, CJ Harper, and J Nembrini. Towards dependable swarms and a new discipline of swarm engineering. In Erol Sahin and William Spears, editors, Swarm Robotics Workshop: State-of-the-art Survey, number 3342, pages 126–142, Berlin Heidelberg, 2005. Springer-Verlag.
  - Winfield AFT and Nembrini J, 'Safety in Numbers: Fault Tolerance in Robot Swarms', Int. J. Modelling Identification and Control, 1 (1), 30-37, 2006.
  - Winfield AFT, Liu W, Nembrini J and Martinoli A, 'Modelling a Wireless Connected Swarm of Mobile Robots', Swarm Intelligence, 2 (2-4), 241-266, 2008.
  - AFT Winfield, 'Foraging Robots', in Encyclopedia of Complexity and Systems Science, Editor-in-chief: Robert A Meyers, Springer, 2009.
  - AFT Winfield and F Griffiths, 'Towards the Emergence of Artificial Culture in Collective Robot Systems', in Symbiotic Multi-robot Organisms, Eds. P Levi and S Kernbach, Springer, 2010.
  - Bjerknes JD and Winfield AFT, 'On Fault-tolerance and Scalability of Swarm Robotic Systems', in Proc. Distributed Autonomous Robotic Systems (DARS 2010), Lausanne, November 2010.
  - Lachlan Murray, Wenguo Liu, Alan Winfield, Jon Timmis, and Andy Tyrrell, Analysing the Reliability of a Self-reconfigurable Modular Robotic System, in Proc. 2011 International ICST Conference on Bio-Inspired Models of Network, Information and Computing Systems (BIONETICS 2011), York, December 2011.
  - Dixon C, Winfield A and Fisher M (2011), Towards Temporal Verification of Emergent Behaviours in Swarm Robotic Systems, in Proc. Towards Autonomous Robotic Systems (TAROS 2011), Sheffield, September 2011.
  - Liu W and Winfield AFT, Modelling and Optimisation of Adaptive Foraging in Swarm Robotic Systems, International Journal of Robotics Research, 29 (14), 2010.
  - Bjerknes JD, Winfield AFT and Melhuish C, 'An Analysis of Emergent Taxis in a Wireless Connected Swarm of Mobile Robots', Proc. IEEE Swarm Intelligence