



Swarm Engineering: a bio-inspired approach to resilient multi-robot systems

Alan FT Winfield

Bristol Robotics Laboratories

<http://www.brl.ac.uk>

Software Engineering for Resilient Systems (SERENE 2010)
Birkbeck College, London
16 April 2010

This Talk



- In three parts:
 - About the Bristol Robotics Laboratory
 - Introduction to Swarm Robotics
 - Potential and Challenges
 - Case Studies in Swarm Robotics
 - Flying flock
 - Symbion project
 - Artificial Culture project
 - Adaptive Swarm Foraging
 - In depth: Wireless connected swarm

About the BRL



- Founded in 1993 as the Intelligent Autonomous Systems Lab
- The Bristol Robotics Lab is, since 2005, a joint research lab of UWE and the University of Bristol

Three main areas:

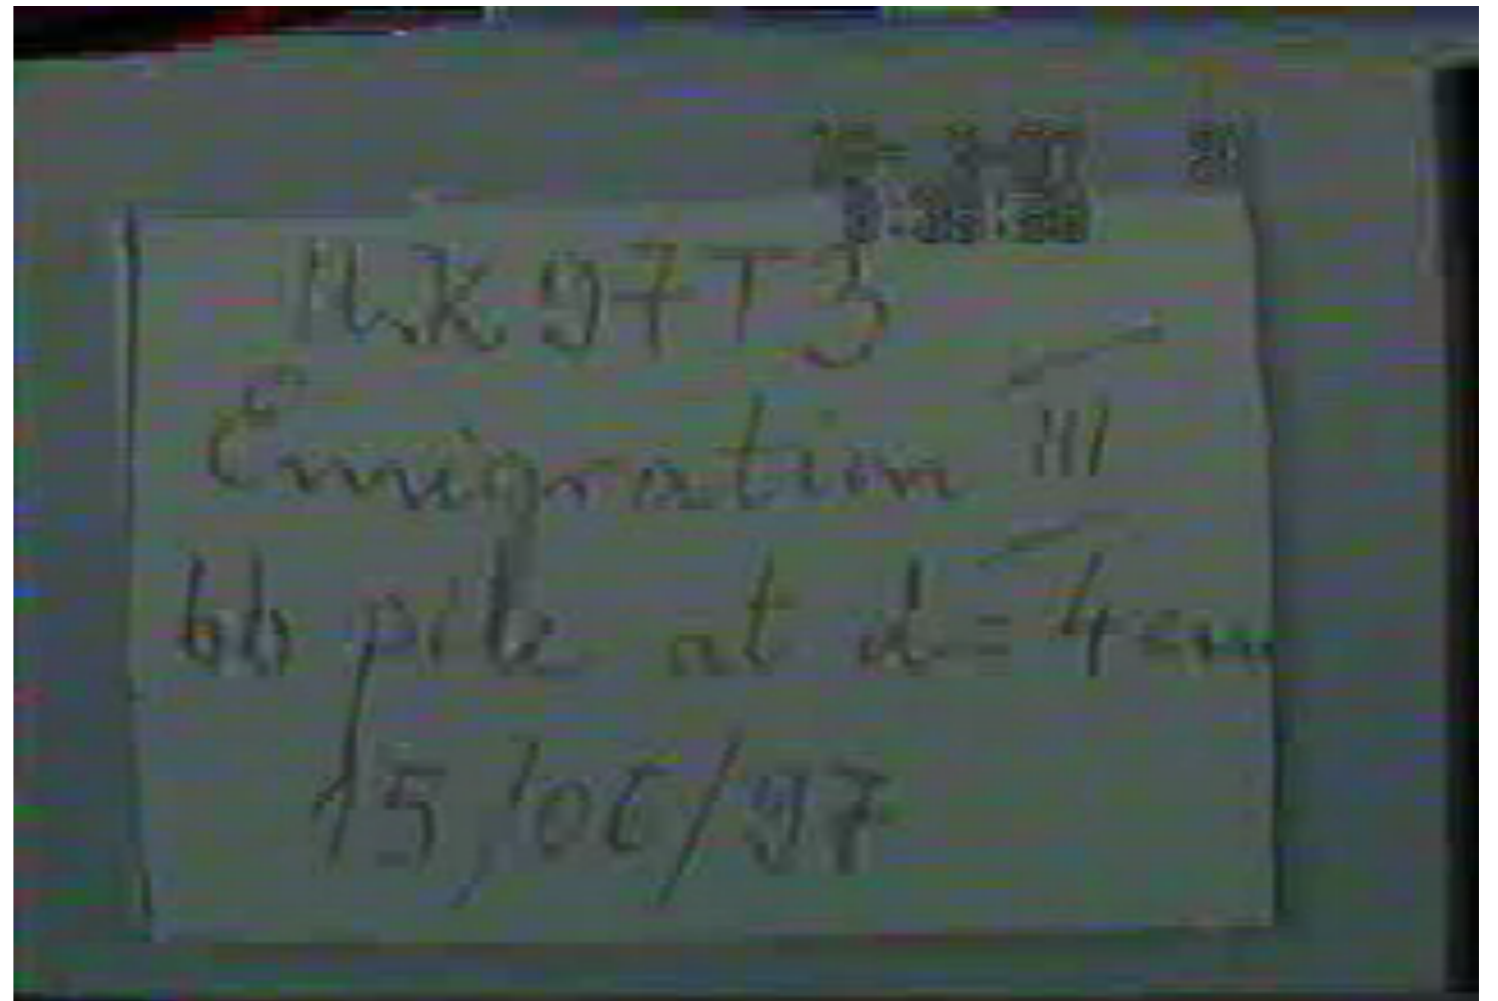
- Biological Robotics
- Human Robot Interaction
- Swarm Robotics

Strong focus on *real robot experiments*



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



Lepto thorax 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..?

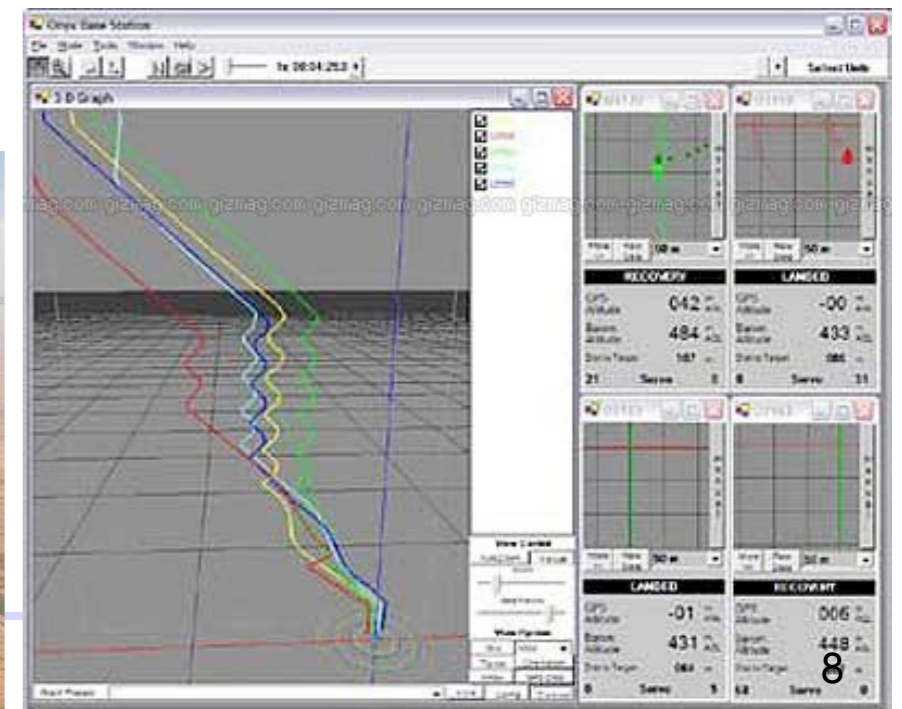
- What are the design principles involved?
 - how do we determine the *local rules* for each individual agent, in a principled way?
- How can we validate overall behaviours that are *emergent* properties?
 - notwithstanding these (difficult) questions...
- A powerful new engineering paradigm for large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

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/>



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]

Symbion


A 5 year project to build a Symbiotic Evolutionary Robot Organism

Individual robots are, in effect, 'cells' in a multi-celled organism, which

- ✓ self-assemble
- ✓ differentiate
- ✓ share resources
- ✓ evolve and adapt

The Symbion organism could have, for example...

- ✓ homeostasis
- ✓ an auto-immune system




Jasmine
robots



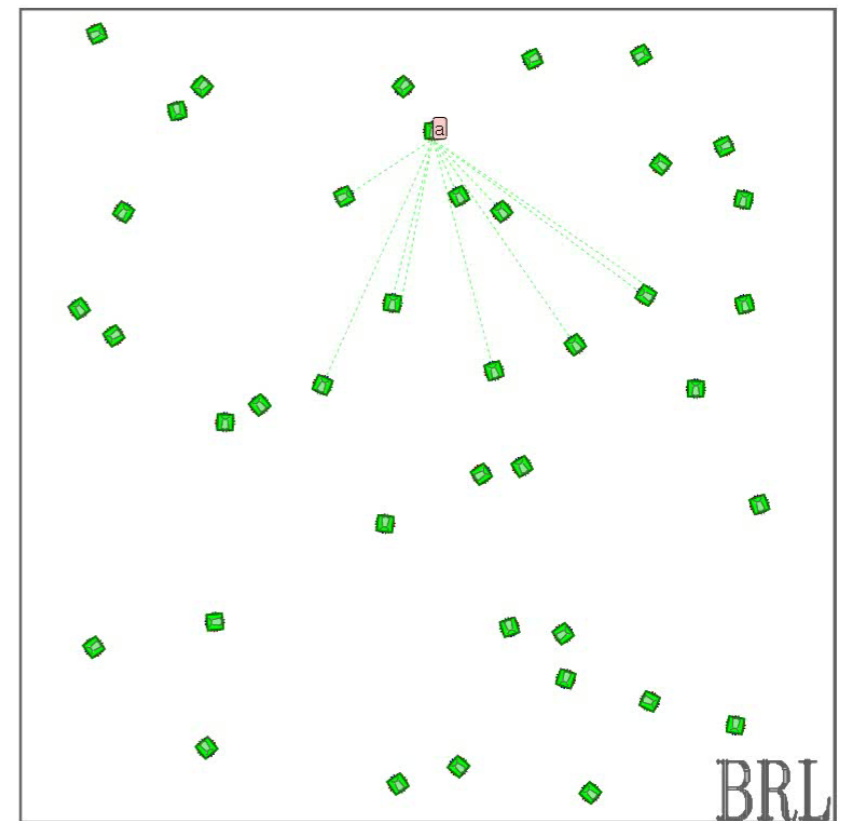


3D physics simulation of artificial organism (Karlsruhe)

Development of swarm to organism
2D morphogenesis (BRL)



Symbion robots - April 2010




0m 17s 500msec [4.9]

The Emergence of Artificial Culture in Robot Societies

Social learning in collective robotics

- our aim is to model the *processes* and *mechanisms* of the emergence of culture in social agents...
- by introducing robot-robot imitation (social learning) to model and study the propagation of robot *memes* across the robot society



e-Puck robots



Artificial Culture Lab

A minimalist wireless connected swarm


- Research question: is it possible to maintain swarm integrity (aggregation) using wireless alone?
- In other words:
 - Is it possible to use wireless connectivity as a *structural component* in building swarm systems..?
- We seek simple rules linking locomotion with communications
 - To create emergent swarm coherence and
 - Scalable control of swarm morphology

A Minimalist Approach


- Robots have
 - Range limited, omni-directional wireless communications
 - *Situated communications*
 - Robots can transmit their identity, but signal strength not available
 - No global positional information
 - No range or bearing sensors
 - Only local knowledge of connectivity

Primitive behaviour


(i) Connected




(ii) Connection lost




(iii) Turn back



(iv) Reconnected, choose new random heading

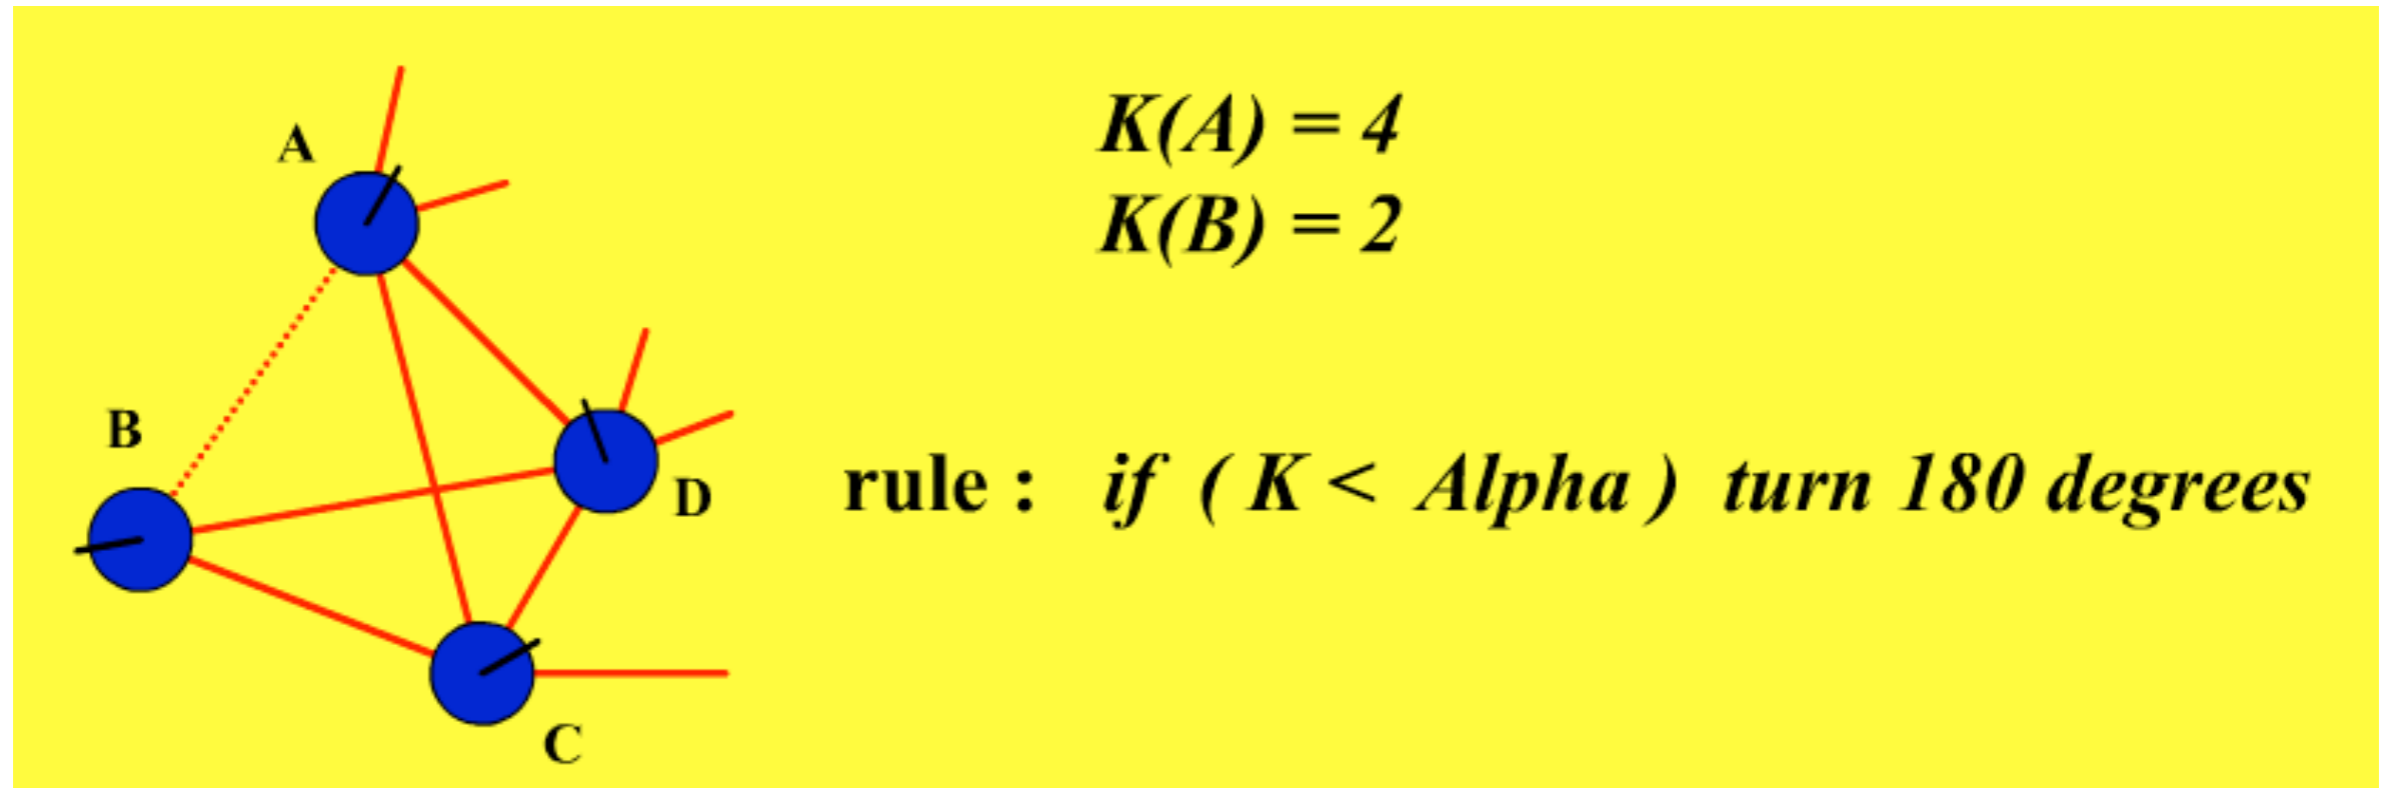


Primitive behaviour
running on 2
Linuxbots





Basic Algorithm

- Extend the basic primitive to multiple robots...
 - React to the number of neighbours in range, i.e. the number of connections K





Physical Implementation

- Experimental platform: the LinuxBot*



Emergent ad-hoc wirelessly connected network


- Single parameter area control:
 - swarm disposition for $\alpha = 5$, $\alpha = 10$



Mathematical Modelling

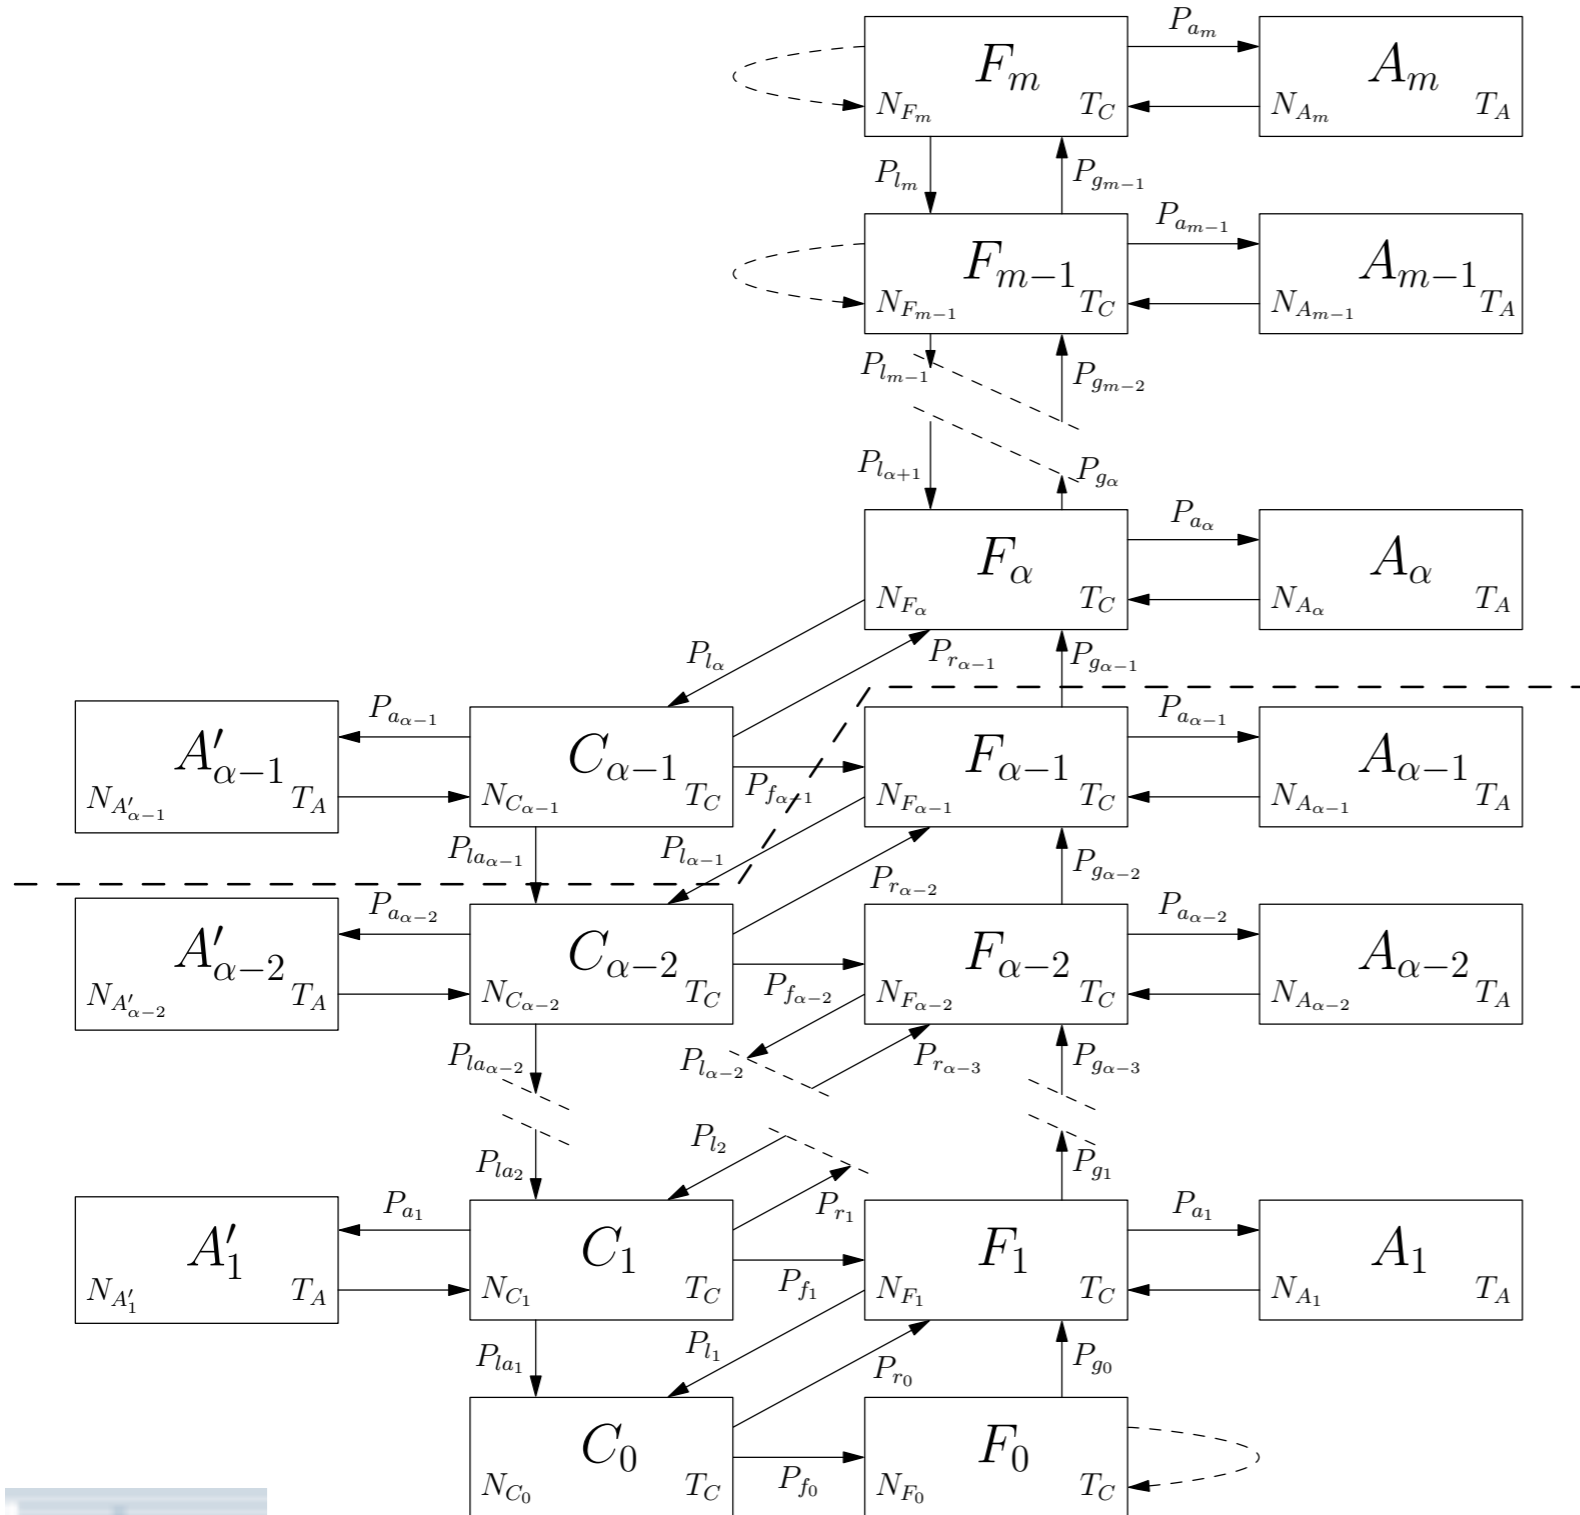
- We model the wireless connected swarm, by extending the probabilistic approach of Martinoli *et al**.
- We take the Finite State Machine (FSM)
 - express as an ensemble of probabilistic FSMs...which lead to a set of difference equations
 - geometrically estimate the transition probabilities
 - compare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

Simplified Finite State Machine



- Avoidance behaviour: triggered by short-range collision sensor
- Coherence behaviour: triggered by number of wireless connections falling below the threshold α

Probabilistic PFSM



Each box represents the number of robots in the swarm:

- *in a given state, and*
- *with a given number of connections*

The PFSM thus describes the state/connection structure of the swarm.

Transition Probabilities

probabilities	comments
P_a	collision with another robot
P_l	loss of a connection in <i>forward</i> state
P_g	gain of a connection
P_r	recovery of a connection
P_f	failure to recover a connection
P_{la}	loss of a connection in <i>coherence</i> state

Difference Equations

- We can now write expressions for the change in number of robots in each state from one time step to the next
 - for the *avoidance* state

$$N_{A_i}(k+1) = N_{A_i}(k) + P_{a_i}N_{F_i}(k) - \Delta_{A_i}(k+1 - T_A)$$

$$\Delta_{A_i}(K+1) = P_{a_i}\Delta_{F_i}(k+1 - T_A)$$

- and so on for other states

- There are N robots in the swarm, so

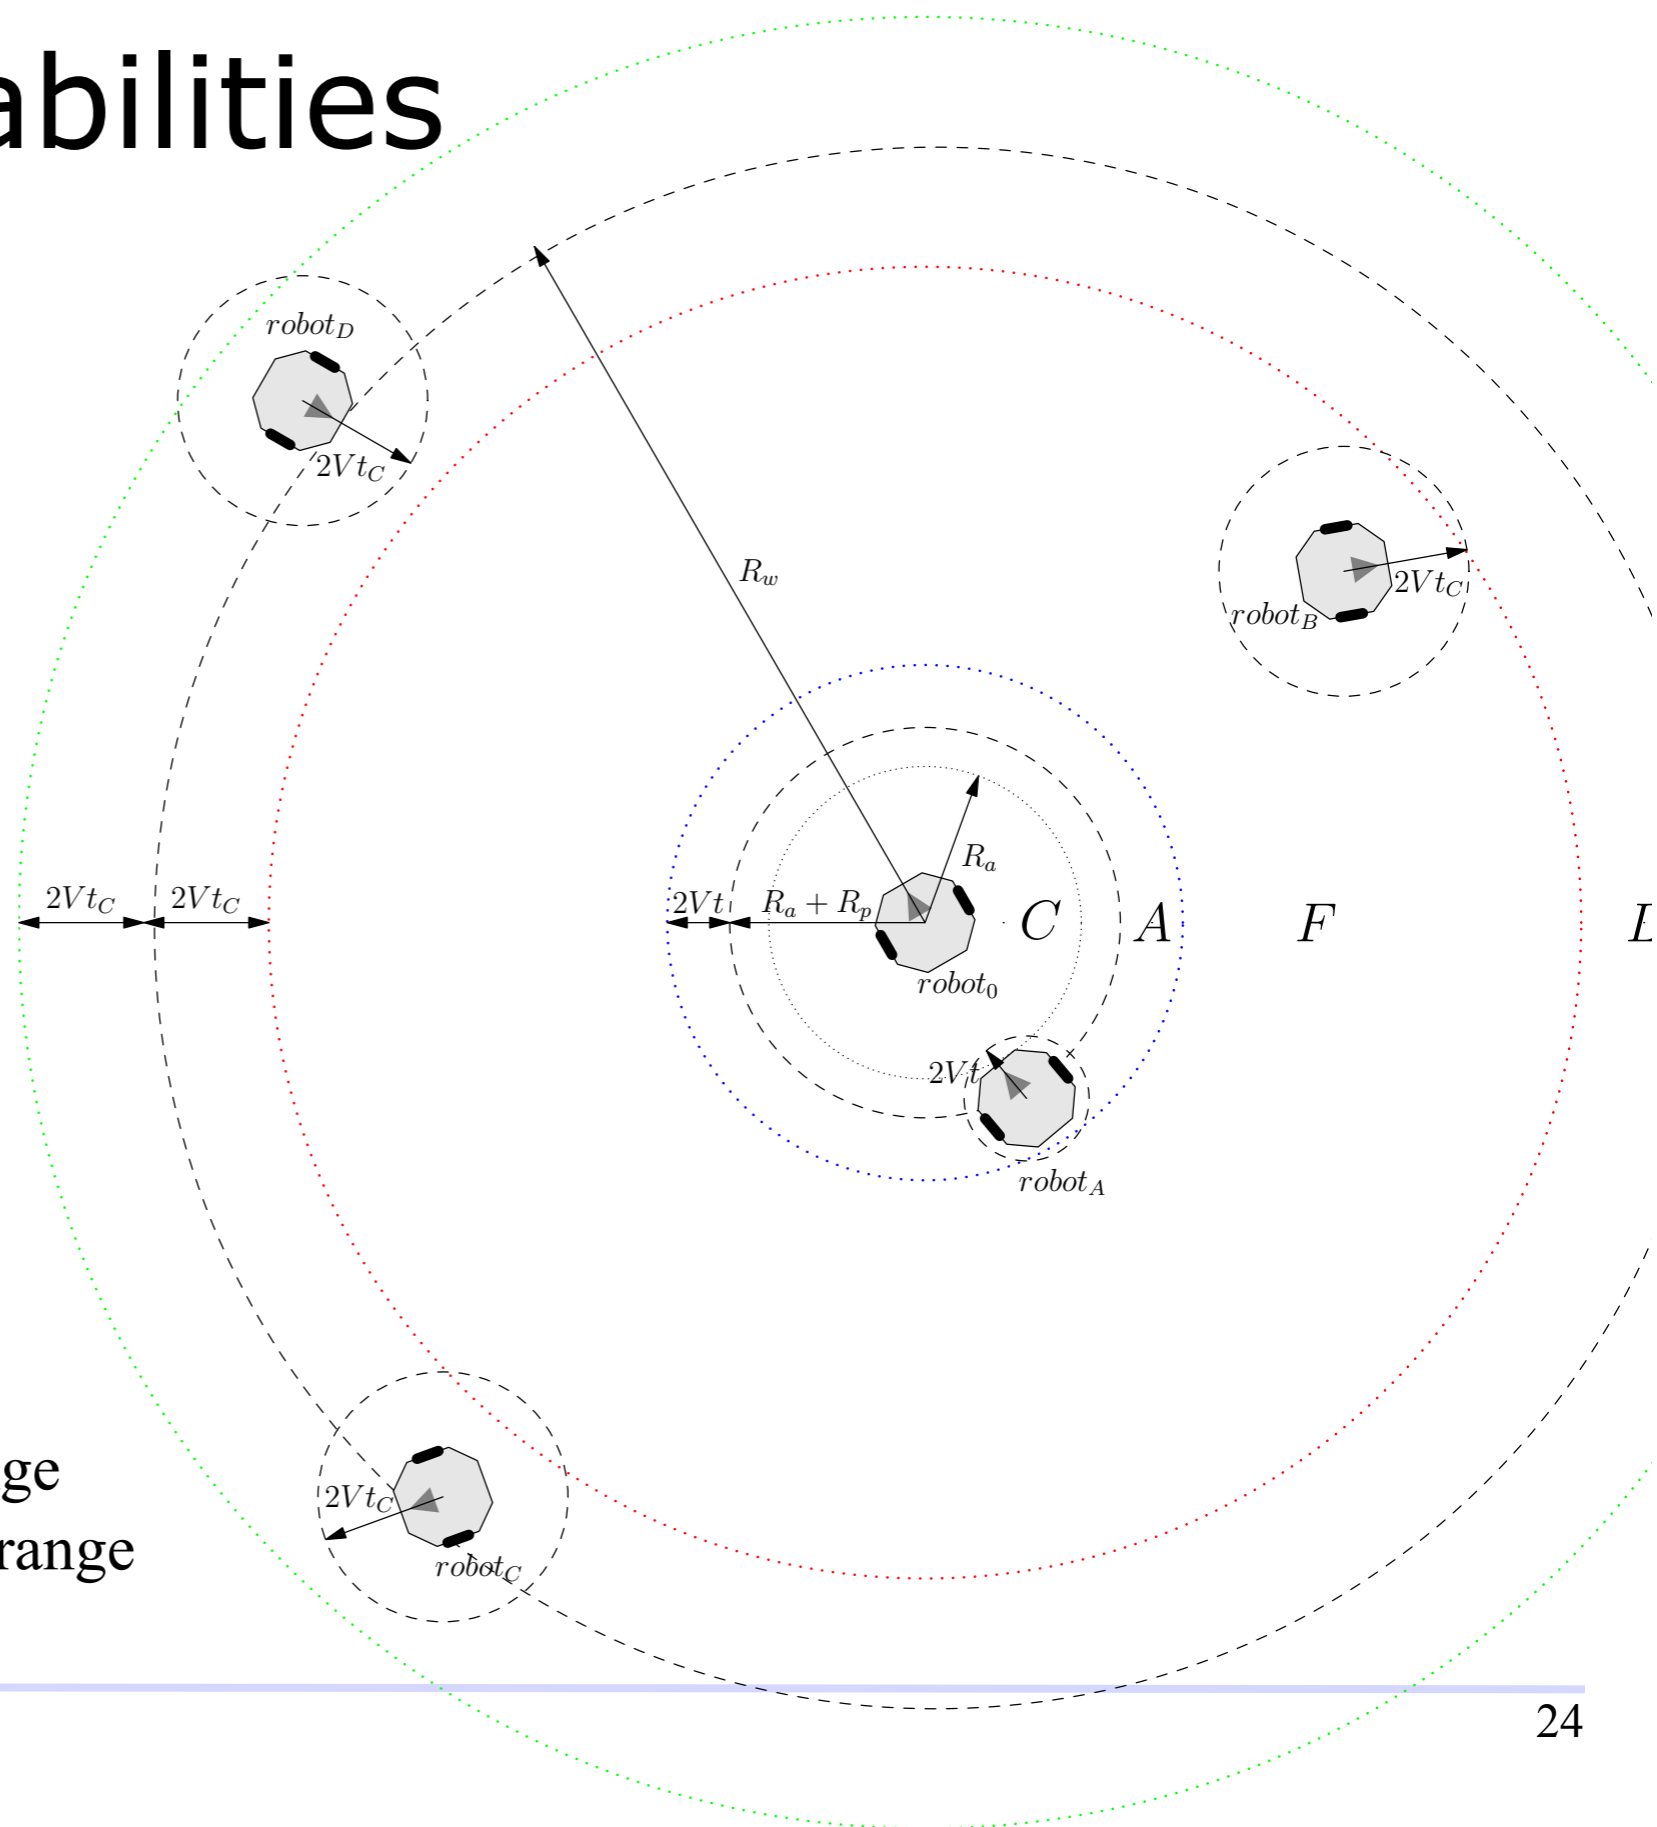
$$N = \sum_{i=1}^{\alpha-1} N_{A'_i}(k) + \sum_{i=1}^m N_{A_i}(k) + \sum_{i=0}^m N_{F_i}(k) + \sum_{i=0}^{\alpha-1} N_{C_i}(k)$$

Geometrical estimation of transition probabilities


- With respect to $robot_0$

- Region C: potential collision
- Region A: potential avoidance
- Region L: potential connection loss
- Region R: potential connection recovery

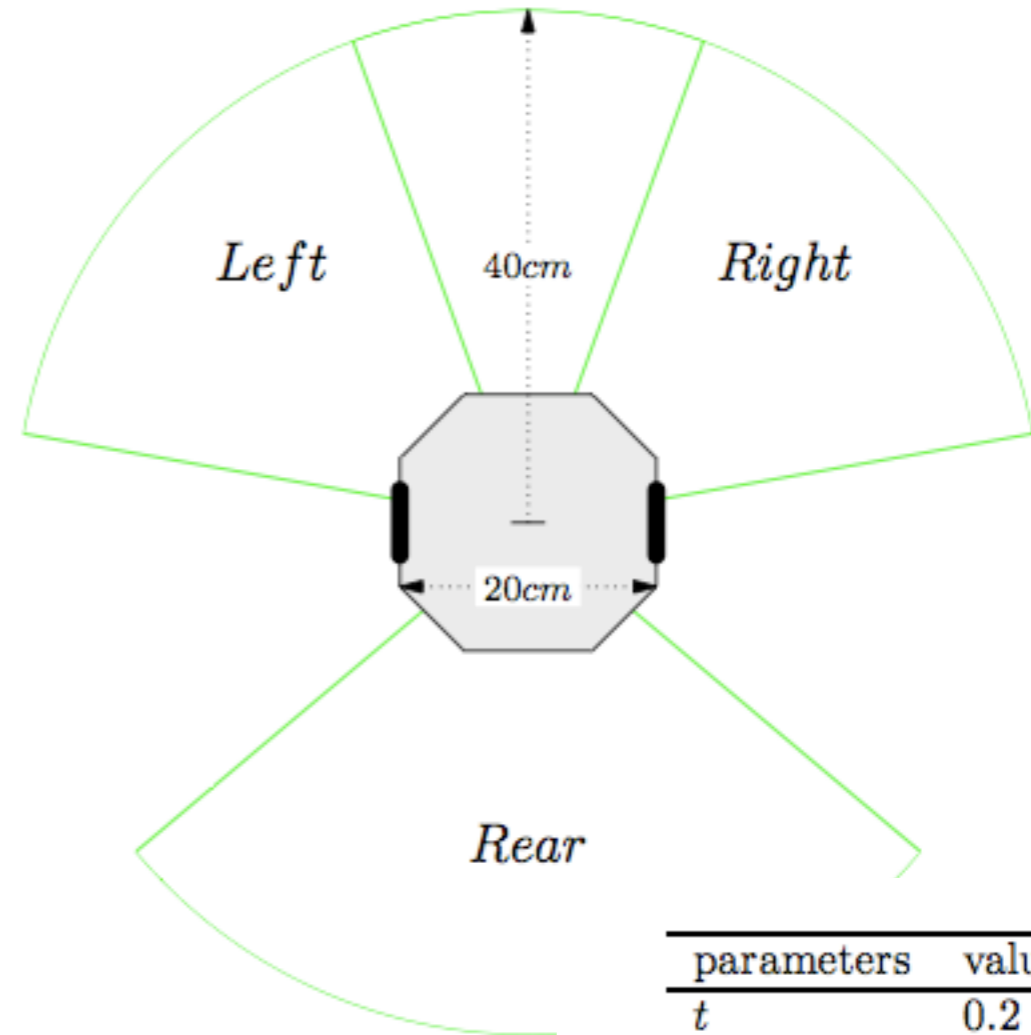
R_a = avoidance sensor range
 R_w = wireless connection range
 V = robot velocity



Simulation for model validation



Blue robots in *forward* state
Red robots in *coherence* state
screenshot from Player/Stage









parameters	values
t	0.2 sec
t_C	3 sec
V	0.15 m/sec
R_a	0.4 m
R_p	0.1 m
R_w	2.0 m

State transition probabilities plotted against connectivity

Left: measured
Right: estimated








Top: $\alpha=5$
middle: $\alpha=10$
bottom: $\alpha=15$



Connectivity: number of robots in states *coherence*, *avoidance* and *forward*

Left: measured
Right: modelled

Top: $\alpha=5$
middle: $\alpha=10$
bottom: $\alpha=15$



Discussion

- We have made a number of simplifying assumptions, primarily
 - in the PFSM we assume connections are lost or gained one-at-a-time
 - in practice more than one connection could be lost or gained in the time T_c between connectivity updates
 - we assume robots uniformly distributed
 - we assume linear functions for $A(x)$, $F(x)$ and $C(x)$
- Despite these assumptions the model achieves excellent qualitative and reasonable quantitative performance

Using Temporal Logic to Specify Emergent Behaviours

- We now investigate the use of a Linear Time Temporal Logic to specify (and possibly prove) emergent properties
- NASA have explored formal methods within the Autonomous Nano-Technology (ANTS) project
 - (Rouff et al, 2004)
 - however that work did not investigate a *temporal logic*

A linear time Temporal Logic

- Extends classical logic with temporal operators,
 - $\bigcirc \varphi$ is satisfied if φ is true in the *next* moment in time
 - $\blacklozenge \varphi$ is satisfied if φ is true at *some* future moment in time
 - $\square \varphi$ is satisfied if φ is true at *all* future moments in time
- Concurrency modelled by interleaving

Specify primitive robot behaviours

Specify the movement primitives, bottom-up

$$\text{moveN}(i) := (\bigcirc x_i = x_i) \wedge (\bigcirc y_i = y_i + a)$$

$$\text{turn180Move}(i) :=$$

$$\begin{aligned} & (\theta_i = S) \wedge (\bigcirc \theta_i = N) \wedge \text{moveN}(i) \vee \\ & (\theta_i = W) \wedge (\bigcirc \theta_i = E) \wedge \text{moveE}(i) \vee \\ & (\theta_i = N) \wedge (\bigcirc \theta_i = S) \wedge \text{moveS}(i) \vee \\ & (\theta_i = E) \wedge (\bigcirc \theta_i = W) \wedge \text{moveW}(i) \end{aligned}$$

One of the four possible state/movement transitions

$$\text{forwardNotConnected}(i) :=$$

$$\begin{aligned} & (\text{motion}_i = \text{forward}) \wedge \neg \text{connected}(i) \wedge \\ & (\bigcirc \text{motion}_i = \text{coherent}) \wedge \text{turn180Move}(i) \end{aligned}$$

Overall swarm specification

Each robot must satisfy both Safety and Liveness properties at all future times

$$Robot_i := \square(Safety_i \wedge Liveness_i)$$

Then specify the Swarm as the logical 'and' of all the robots

$$Swarm := \\ Robot_1 \wedge Robot_2 \wedge \dots \wedge Robot_N \wedge \\ \square(\pi_1 \oplus \pi_2 \oplus \dots \pi_N)$$

Ensure that only 1 robot taking action at a time

Specification of Emergent Properties

Eventually each robot will be connected to at least k distinct others

First specify the emergent properties

Now attempt to prove (or disprove) that the swarm of robots satisfies the emergent behaviours

Each robot is always connected



$$property1 := \Box \Diamond (\forall i \in robotSet. connected(i))$$

$$property2 := \Diamond \Box (\forall i \in robotSet. (\exists j_1 \in robotSet\{i\}. inRange(i, j_1) \wedge \exists j_2 \in robotSet\{i\}. inRange(i, j_2) \wedge \dots \exists j_k \in robotSet\{i\}. inRange(i, j_k) \wedge distinct(j_1, j_2, \dots, j_k)))$$


$$Swarm \Rightarrow property1$$
$$Swarm \Rightarrow property2$$

Reliability Modelling: emergent swarm taxis

- How does it work...
- Robots have simple aggregation:
 - short range: obstacle avoidance (repulsion)
 - longer range: maintain number of connected neighbours (attraction)
- Each robot also has a simple beacon sensor
 - symmetry breaking mechanism: *illuminated* robots have a slightly larger avoid radius than *occluded* robots



Swarm taxis with failures



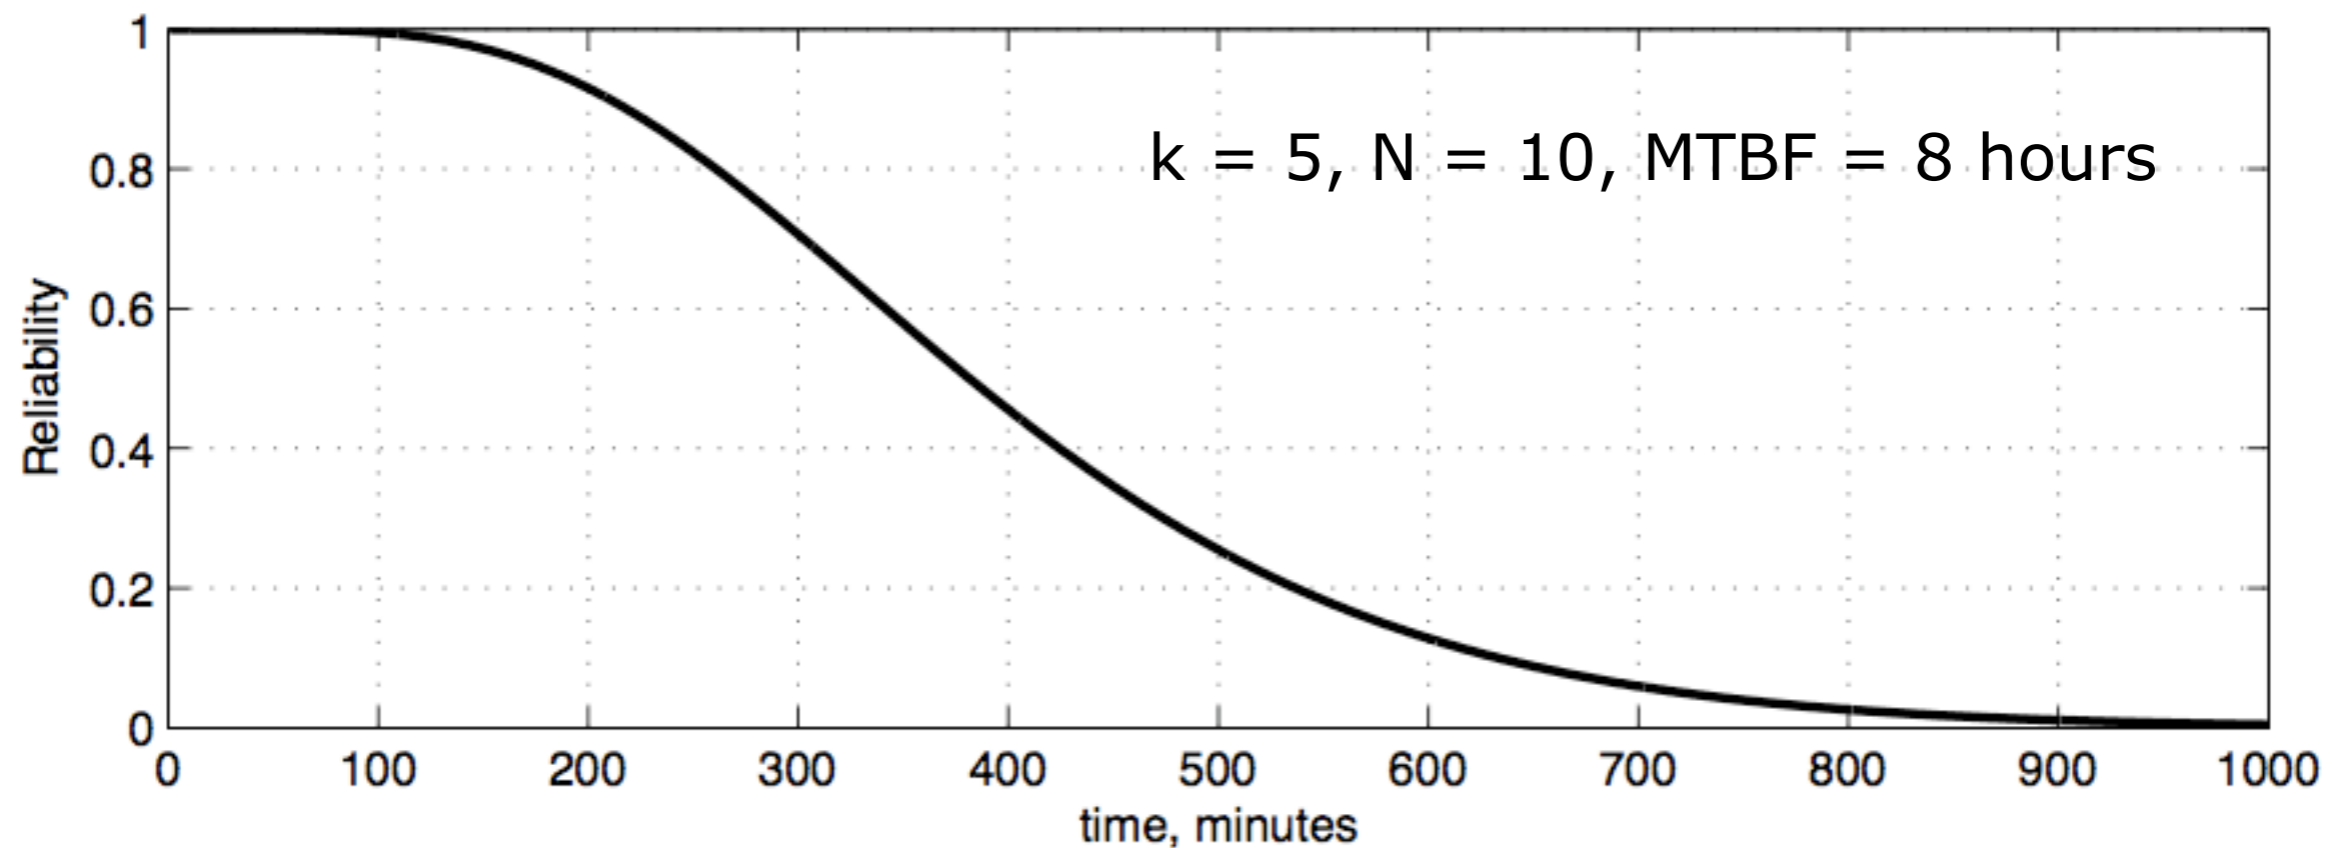
We then introduce worst-case *partial* failures - i.e. robots whose motors fail, but sensing and communications remains ok

Bjerknes 2010


The k-out-of-N reliability model

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


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



Swarm self-repair



Single robot complete failure H5



Single robot partial failure H1

Estimate k for partial failure H1

- Conservatively $k = 0.9N$
 - in other words, we believe the swarm can tolerate 10% of H1 failures at any one time (i.e. within 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. $t = d/2v$, $d = \text{swarm diameter}$, $v = \text{swarm velocity}$

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



Discussion

- The frequent assumption, that swarm systems are automatically scalable and robust, is seriously incorrect
- This result strongly suggests that scaling systems (which rely on emergence or self-organising mechanisms) requires more sophisticated internal mechanisms for dealing with worst-case failures:
 - an *immune system*

Thank you!

EPSRC

Engineering and Physical Sciences
Research Council



- Acknowledgements, colleagues in the BRL, but especially:
 - Dr Chris Harper, Dr Julien Nembrini, Dr Wenguo Liu, Dr Jan Dyre Bjercknes
 - Prof Chris Melhuish, Dr Tony Pipe
- Relevant publications:
 - A. F. T. Winfield, C. J. 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, Sa J, Gago MC, Dixon C and Fisher M, 'On Formal Specification of Emergent Behaviours in Swarm Robotic Systems', *Int. J. Advanced Robotic Systems*, 2 (4), 363-370, 2005.
 - Harper CJ and Winfield AFT, 'A Methodology for Provably Stable Behaviour-based Intelligent Control', *Robotics and Autonomous Systems*, 54 (1), 52-73, 2006.
 - 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', *Encyclopedia of Complexity and Systems Science*, Editor-in-chief: Robert A Meyers, Springer, 2010.