Swarm Intelligence

Rod Goodman
2008 Carnegie Centenary Professor
Edinburgh University
Scotland

27th August 2008

Acknowledgements

This talk describes the work of many good friends and colleagues, in particular:

- *The former students and staff of my Collective Robotics Research Group at Caltech.*
- *Prof. Alcherio Martinoli, Distributed Intelligent Systems and Algorithms Laboratory, École Polytechnique Fédérale de Lausanne, Switzerland.*
- *Prof. Owen Holland, Robotics Group, University of Essex.*
- *Prof. Alan Winfield, Bristol Robotics Group, University of the West of England.*

Swarm Intelligence

Swarm Intelligence: "Any attempt to design algorithms or distributed problem-solving devices inspired by the collective behavior of social insect colonies and other animal societies." [Bonabeau, Dorigo, and Theraulaz, 1999]

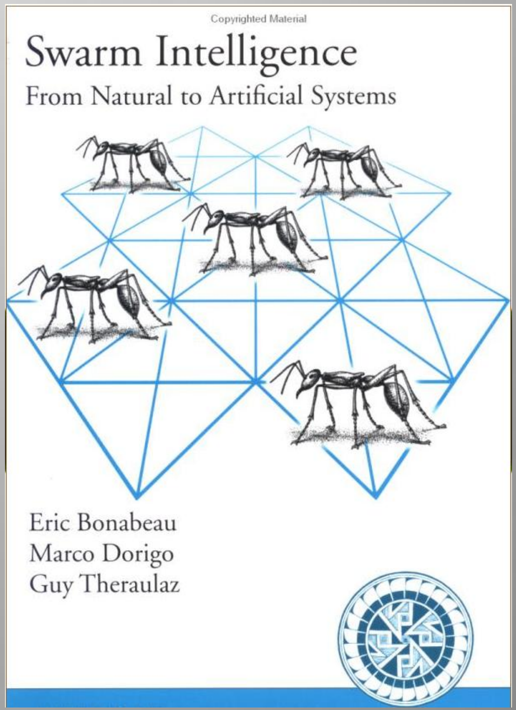

***Biological Inspiration* : social insects (ants, bees, termites) flocks of birds, herds of mammals, schools of fish, packs of wolves, pedestrians, traffic.**

***Emergence:* seemingly intelligent reliable *global* behavior (nest building, foraging, defending the nest, flocking and herding) *emerges* from the collective actions of simple, unintelligent , and unreliable *local*, *distributed*, agents. (**Neurons per ant 300K x 10^6 in colony = $3*10^{11}$, Honeybee 850K x 10^5 = $8.5*10^{10}$, Humans 10^{12})

***Engineering Motivation:* Can we use these principles to achieve intelligent behavior from simple distributed low intelligence robots?**

Caltech EE151: Swarm Intelligence Course, 2000-2005, Instructors: Rod Goodman, Alcherio Martinoli, Owen Holland

Old and New

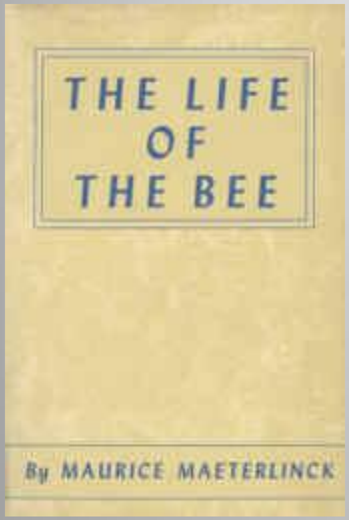

1901

"What is it that governs here? What is it that issues orders, foresees the future, elaborates plans, preserves equilibrium?"

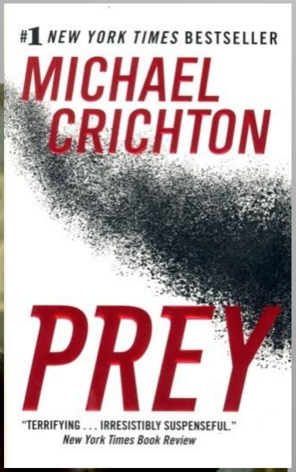

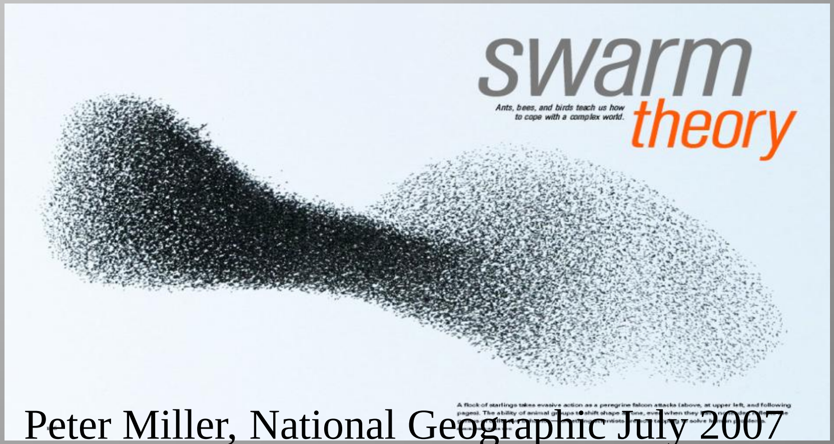

Peter Miller, National Geographic July 2007

The Hope: An Underlying Methodological Framework

Flocks of Birds

Traffic systems

Social insects

Networks of Sensors

?

Schools of Fish

Pedestrians

Multi-robot systems

Outline

- Understanding SI
  - Natural examples
  - Principles
- Modeling SI
  - From Natural Behaviors to Algorithms
- Engineering SI
  - Robotics

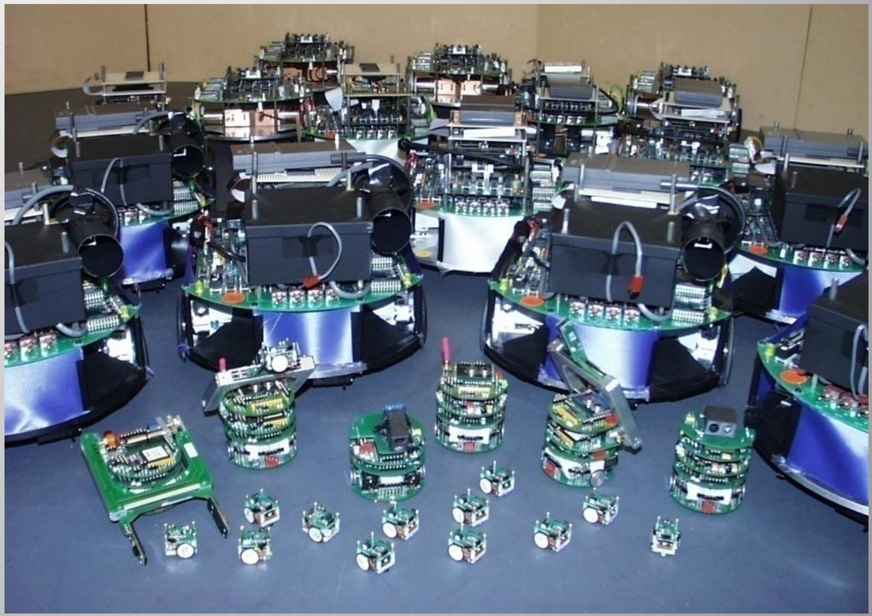

Caltech Collective Robotics Lab Robots: MooreBots, Kephera, Alice

Swarm Intelligence – Natural Examples

Leaf CutterAnts

- Cut and transport leaves, to -
- Feed to a fungus, which they cultivate, which -
- Breaks down toxins in the leaves and produces a mulch rich in sugars, which -
- They eat.
- The first agriculture?

An Ant Bridge

- A lesson in self-sacrifice!

Wasp Nest Construction

A queen starting
construction of a new nest.

Fish Schooling

- Anchovies at the Monterey Aquarium, California

Bird Flight Formation

Pedestrian Swarms

NEW YORK Street Crossing ©DMNikas

Swarm Intelligence-Principles

Characteristics of SI systems

- Intelligent global behavior *emerges* through the collective actions and interactions of many simple individuals.
- Each individual:
  - Follows simple probabilistic behavioral rules, triggered by sensory stimuli
  - Has very limited “intelligence”
  - Has limited local information
  - Self-organizes with no global external control
  - Utilizes Stigmergy as a means of indirect communication via the environment
  - Utilizes Stigmergy to store “state” in the environment

The Ingredients of Self - Organization

- Positive feedback – results in growth, choice, reinforcement
  - e.g. recruitment to a food source : pheromone trails in ants, waggle dance in bees.
  - e.g. attraction to a task: pheromone build-up from dead ants triggers nest cleaning.
- Negative feedback - counteracts and stabilizes positive feedback
  - e.g. pheromone evaporation on trails
  - e.g. saturation, exhaustion, or competition in foraging for food
- Randomness – results in new discovery, adaption
  - Amplification of fluctuations: random walks, errors, random task-switching.
  - e.g. “lost foragers” discover new food source, start recruiting.
- Multiple interactions
  - Allows robust global structure to emerge through self-organization and the reinforcement of many probabilistic individual actions.

Stigmergy

Grassé P. P., 1959

- Was introduced by Pierre-Paul Grasse in the 1950's to describe the indirect communication taking place among individuals in social insect (termite) societies.
- "La coordination des taches, ……
  The coordination of tasks and the regulation of constructions does not depend directly on the workers, but on the constructions themselves. ***The worker does not direct his work, but is guided by it***.
- Hence: **STIGMERGY** (*stigma,* prick, sting, mark; *ergon,* work, product of labor = hence stimulating product of labor)
- Stigmergy occurs when an insect's actions are determined or influenced by the consequences of another insect's previous actions.

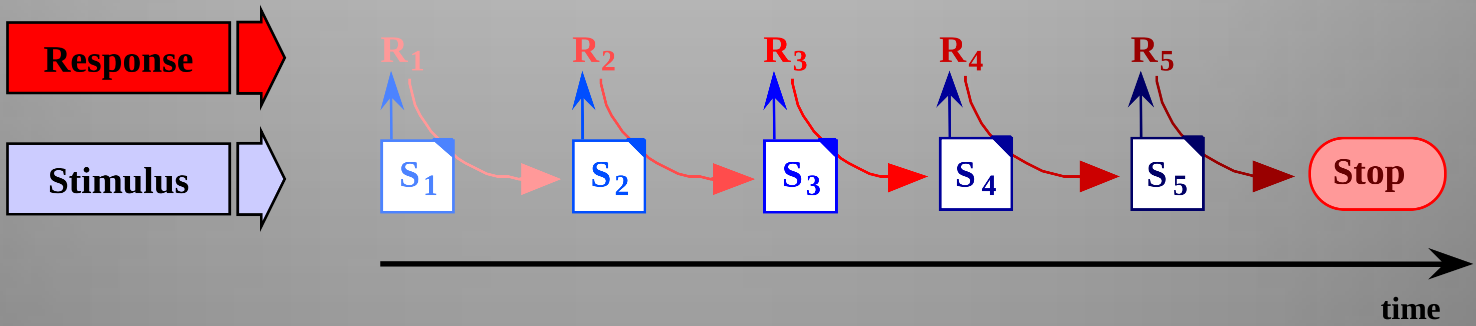

Quantitative or Continuous Stigmergy

- The stimulus is a continuous variable
- The effect is to modulate the action, or to switch to another action
- E.g. Termite Queen Chamber construction: a template, then stigmergy.

•Temperature is critical for reproduction (built in?)
•The pheremones around the queen provide a template for the construction.
•Workers initially build pillars guided by the template
•Pillars are built because the "cement pheremone" attracts more workers to the pillar sites (positive feedback)
•Finally the walls are filled in with sufficient airflow to maintain temperature
•Full circle

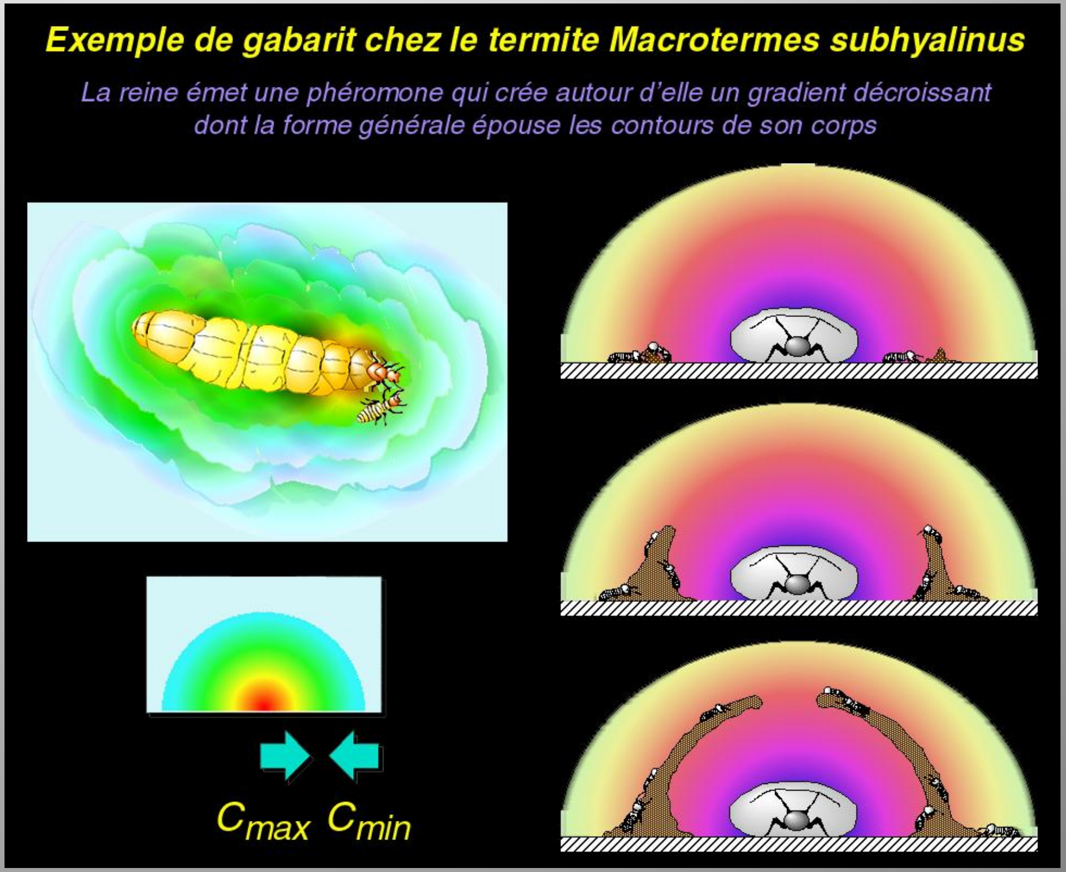

Qualitative or Discrete Stigmergy

- stimulus is a discrete variable
- The effect is usually to switch to another action
- E.G. paper wasp nest construction, an initial template :
  - attach initial stalk like pedicel to substrate (branch, wall)
  - Build two cells on opposite sides
  - Wasps follow rules of where to build the next cell
  - Probability of creating a new cell given the current state of neighboring cells changes due to stigmergy

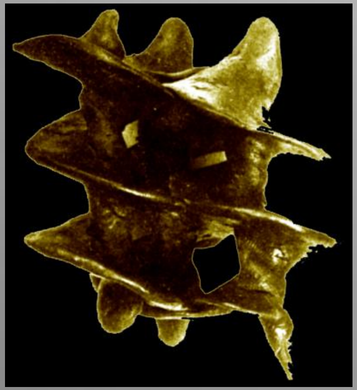

Rules evolve and are modulated by temperature, air flow

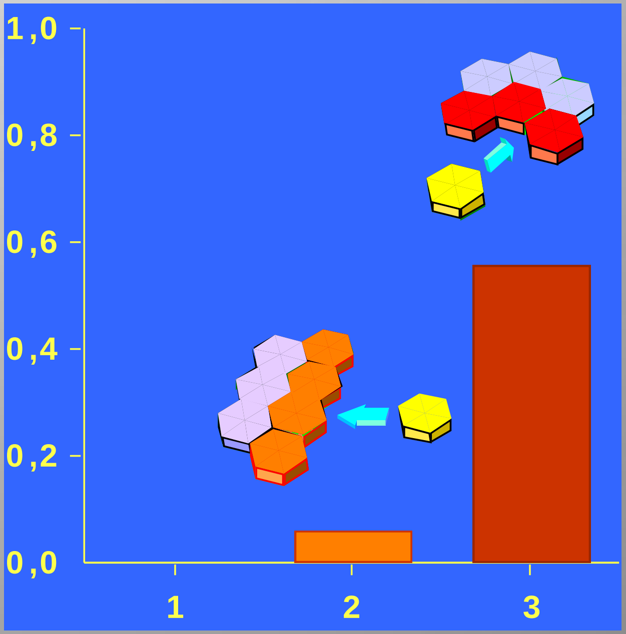

Number of adjacent cell walls

Artificial Construction

- Swarm of nest builders
- Probabilistic rules with local perception

•Stable, repeatable, architectures
•An architecture is stable when repeated simulations with the same rule set generate architectures with the same global structure.

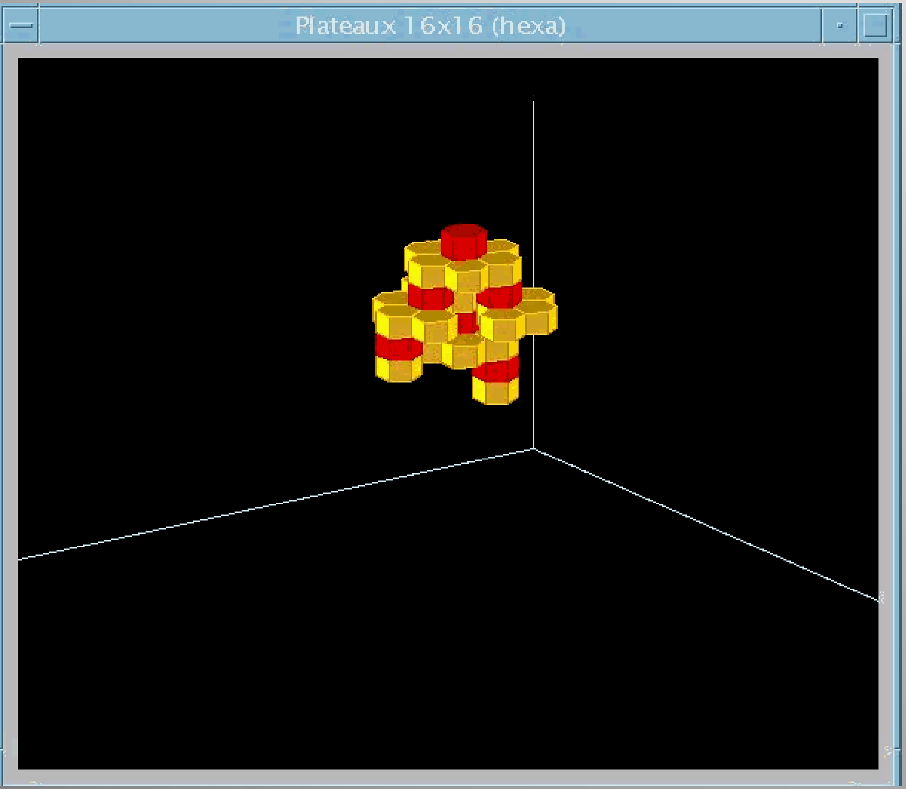

Chatergus (39 rules)

Artificial Nest Structure (35 rules)

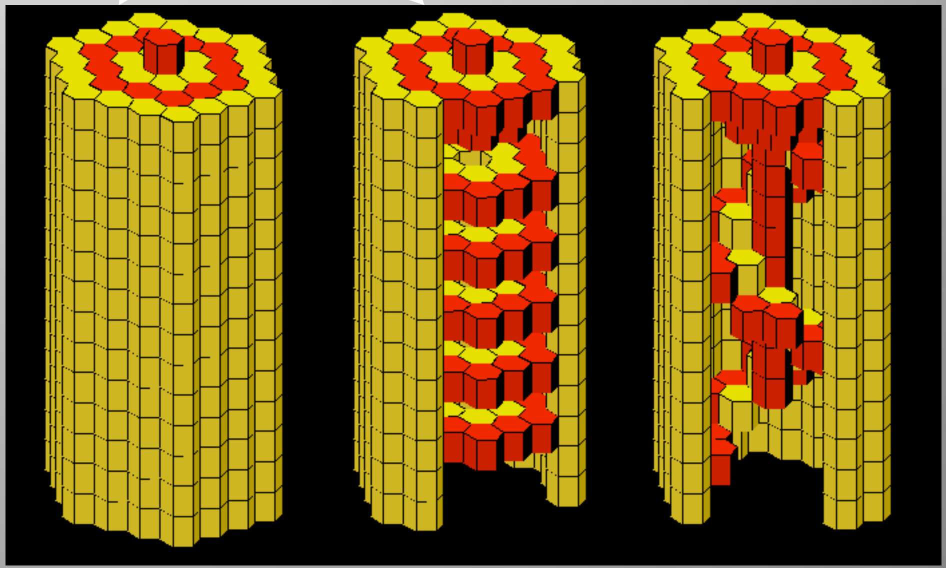

•Can reverse engineer microrules using Genetic Algotithms

Robotic Self-assembly

- Lionel Penrose (1898 – 1972)
  - Self-assembly mechanisms of genetic relevant molecules reproduced with passive wood bricks
  - External energy source (shaking, human action)
  - Evolution and self-assembly tightly coupled.

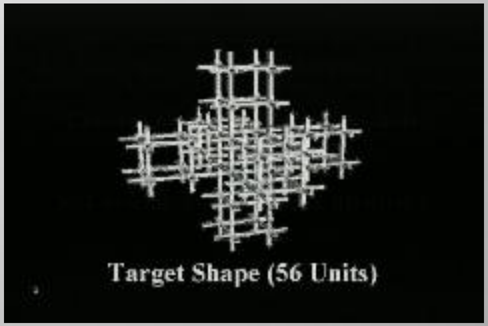

Prof Alan Winfield
Bristol Robotics Lab
UWE

2008 EU Project Concept Video

Modeling Swarm Intelligence – From natural behaviors to algorithms

Ant Navigation

- Different Modalities:
  - Vision: compound eyes, some species have good vision, some have poor vision, some are blind.
  - Detection of polarized sunlight.
  - Perhaps magnetic cues.
  - Mass recruitment to source by chemical trail.
- Pheromone trail laying:
  - Can vary type, amount, and frequency of deposit.
- Pheromone trail following:
  - Sensing via antenna.
  - Strategy: turn towards side with strongest (osmotropotaxis)
  - Ants often move faster on stronger trails.
  - Angle of trail bifurcations (60 degrees) give direction to/from nest.

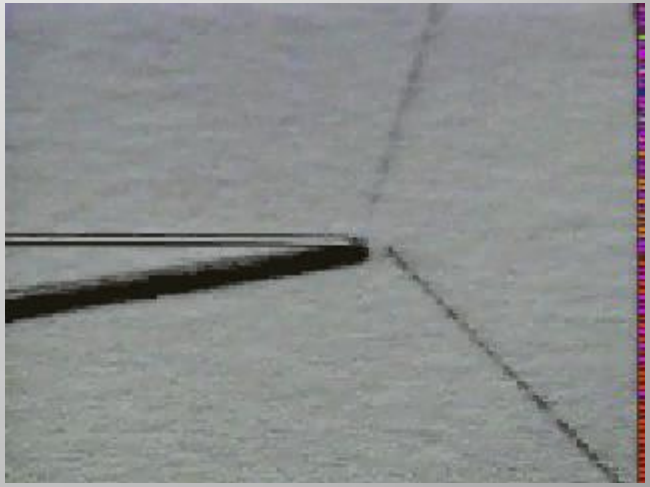

Termite Following a Pheromone Trace

Ant Foraging

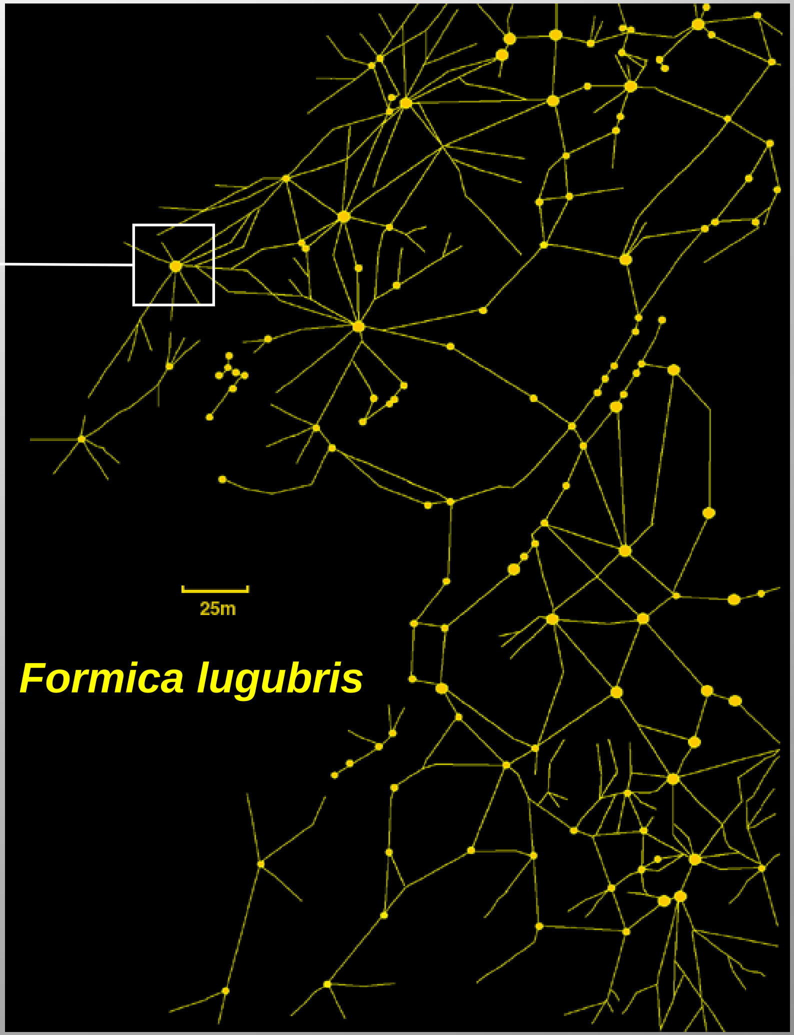

Trail network in a super-colony

Army Ant Raid Fronts

- 500,000 to 20,000,000 per colony.
- Feed on small social insects, arthropods, mixed.
- 20m raid front, 30,000 prey items brought back.
- Bivouac with 50,000 larvae moved every night.
- 15 nomadic days, 20 static days, 14 raids, each at 123 degrees to previous.
- Different species have different raid patterns, depends on environment and prey.
- Perhaps near optimal “distribution network” that maximizes the amount of food brought back to the nest for a given energy expense.

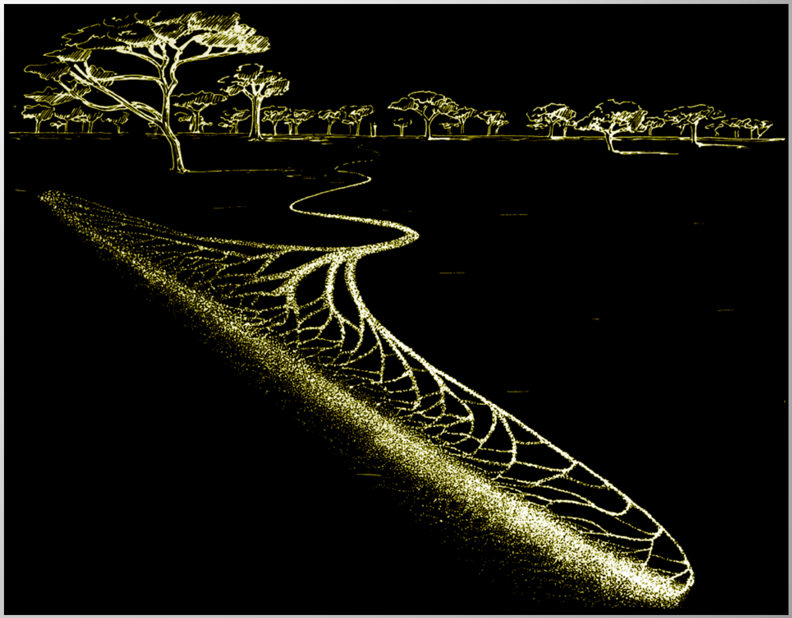

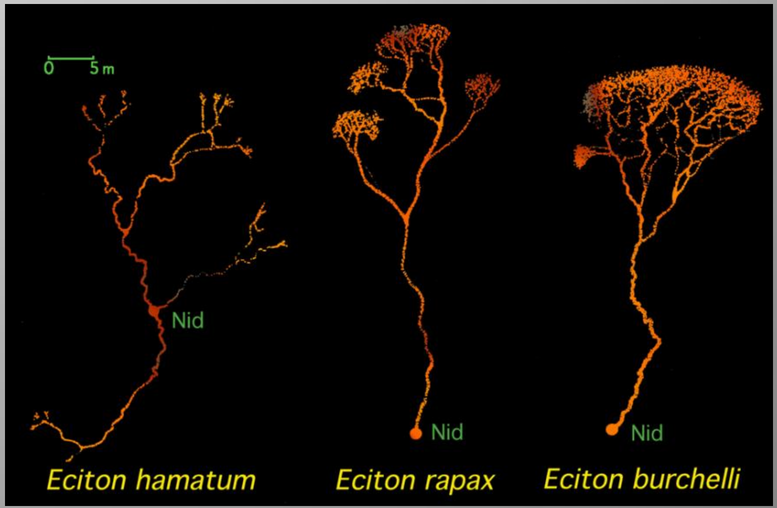

Key Experiment: Suspended, Symmetric Bridge

- Two branches (A and B) of the same length connect nest and food source .
- Positive feedback results in self-organization to one path only.

$$P_A = \frac{(k + A_i)^n}{(k + A_i)^n + (k + B_i)^n} = 1 - P_B$$

A_i : number of ants having chosen branch A
B_i : number of ants having chosen branch B
n : degree of nonlinearity, (n higher = faster triggering for one of the branches)
k : degree of attraction of a unmarked branch, (k greater = greater amount of pheromone needed to make a non-random choice).

Food source

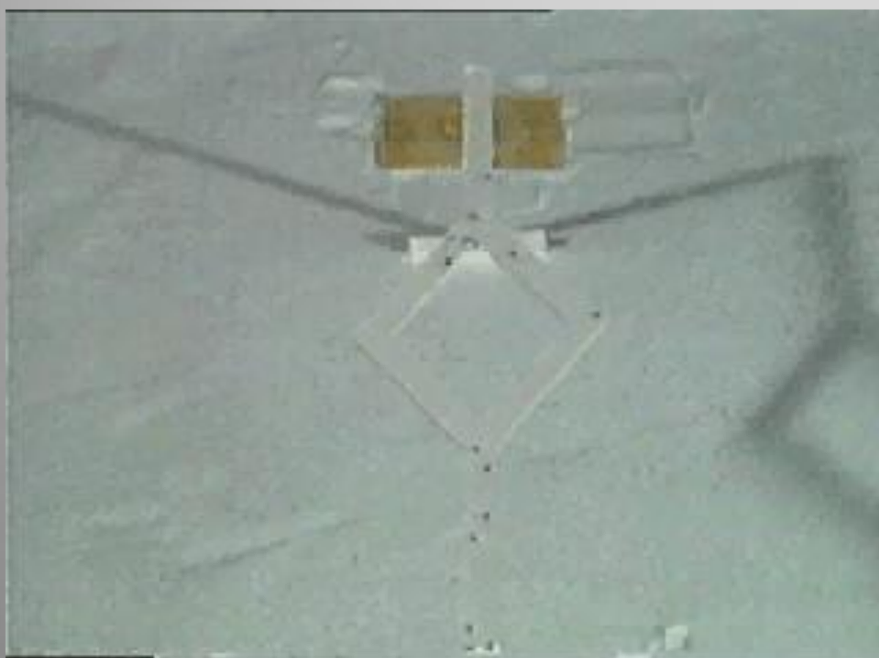

Nest

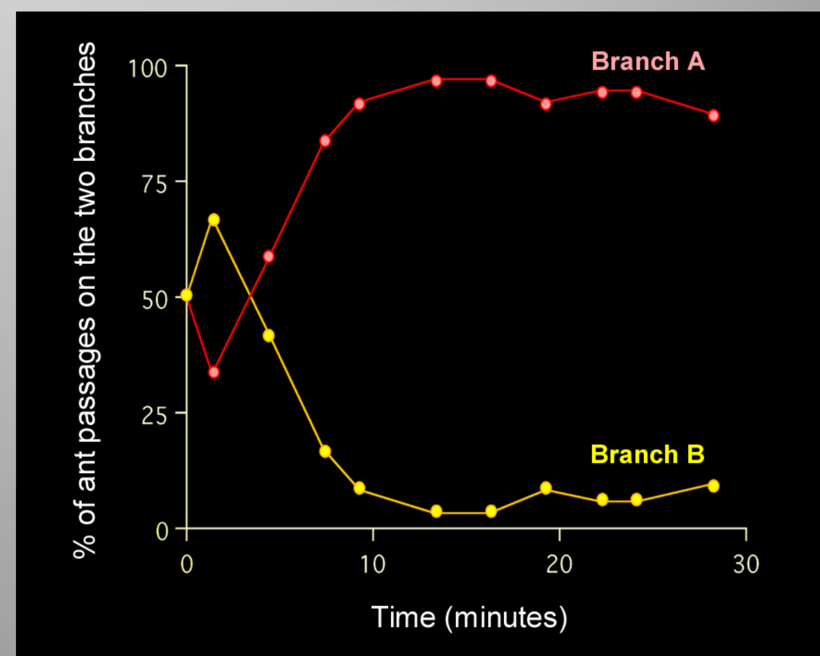

The Suspended, Asymmetric Bridge Experiment

- Two branches (A and B) differing in their length connect nest and food source.
- Test for the optimization capabilities of ants.

Experiments show that the chance of the shorter path being eventually selected increases with the length ratio ***r*** of the two branches.

Food source

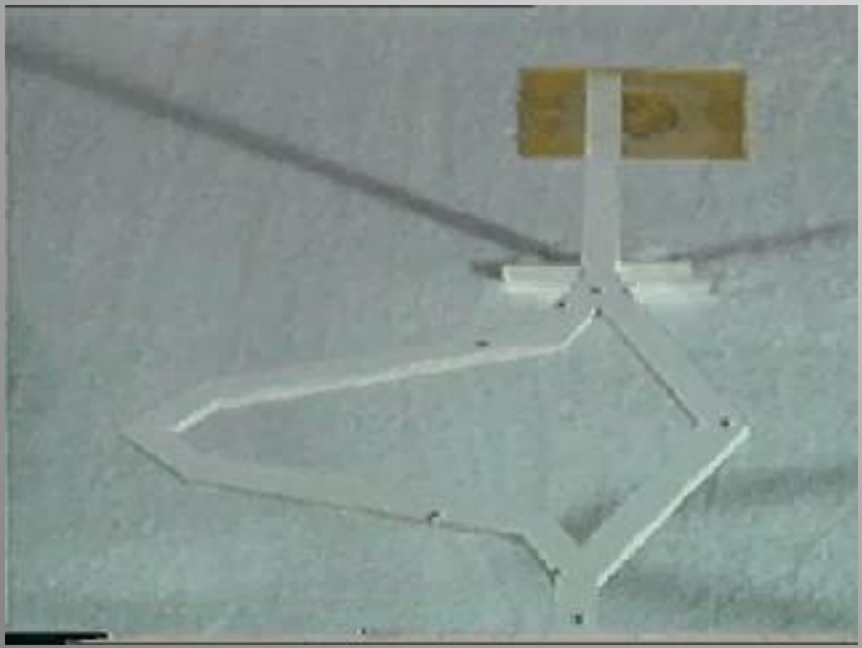

Nest

- If the shorter branch is presented 30 minutes after the longer branch:
- Argentine Ants (*Linepithema Humile*) get stuck on the longer branch because they use mainly pheromone-based navigation.
- *Lasius Niger* ants find the shorter branch because they can sense direction change and hence turn round.

Ant-based Routing in Telecom Networks

- Circuit Switched: “**Ant-based load balancing in telecommunications networks”,** Schoonderwoerd, Holland, Bruten, and Rothkrantz (1996)
- Packet Switched: “**AntNet: A Mobile Agents Approach to Adaptive Routing”,** Gianni Di Caro, Marco Dorigo Universite Libre de Bruxelles, Belgium (1997) .
- Mobile Networks

•ANT RFID Networked Active Tags, InfinID Inc.

- •Ant like Constrained Flooding Algorithm. 2007 ©
- •Reinforcement learning of cost to reach gateway.
- •Probabilistic retransmission based on message (pheromone) frequency.
- •Differential retransmission delays based on cost differences.

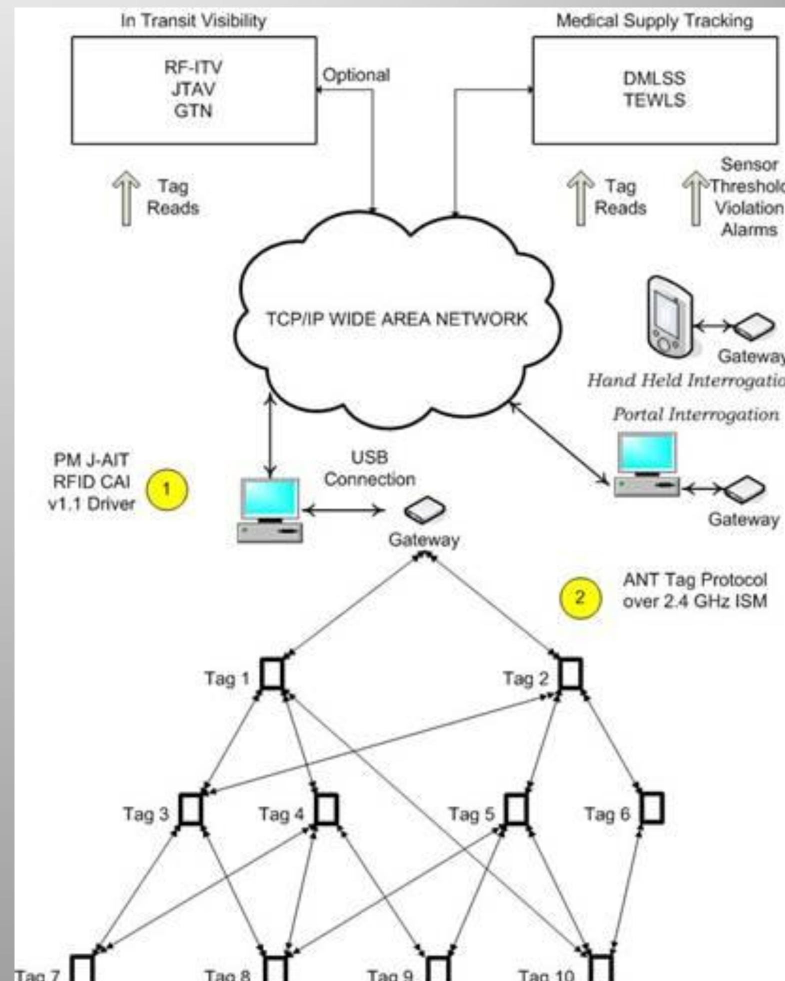

Ant Colony Optimization Algorithms

- Algorithms based on pheremone trail laying have been applied to many combinatorial optimization problems.
- Produce good, but not necessarily best results.
- Best results are usually obtained with specialized algorithms.
- However ACO provides a *general* robust solution.
- Recent developments in Particle Swarm Optimization (PSO) algorithms produce very good results. e.g. Emergency Evacuation Optimizaton:

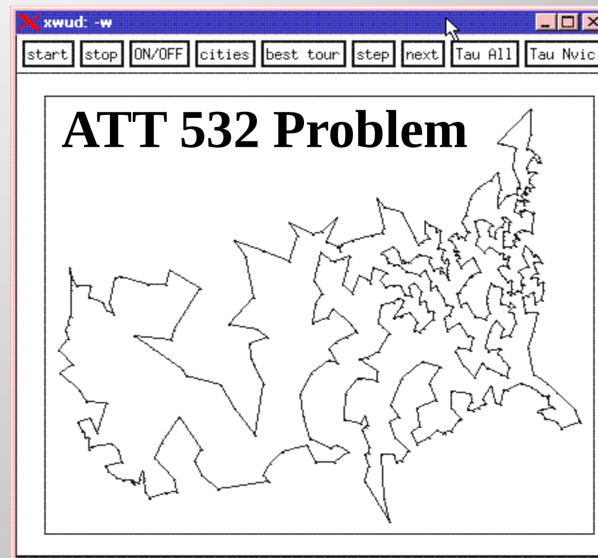

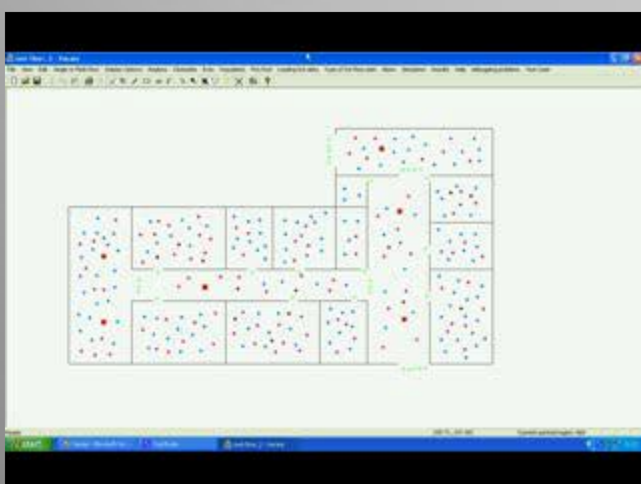

Available applications and results obtained

TRAVELLING SALESMAN PROBLEM (TSP)
Same quality as best GAs, worse than world-champion Iterated Lin-Kernighan

QUADRATIC ASSIGNMENT PROBLEM (QAP)
Best heuristic currently available on structured problems

JOB-SHOP SCHEDULING PROBLEM (JSP)
Promising preliminary results on the single machine total tardiness problem

VEHICLE ROUTING PROBLEM (VRP)
Among the best methods for vehicle routing problems with time windows

SEQUENTIAL ORDERING PROBLEM (SOP)
Best heuristic currently available

GRAPH COLOURING PROBLEM (GCP)
Good results, but not the best

SHORTEST COMMON SUPERSEQUENCE PROBLEM (SCS)
Very good results, ***among the best methods***

MULTIPLE KNAPSACK PROBLEM (MKP)
Promising preliminary results

Division of Labor and Task Allocation

- Reproductive – queens and drones.
- Physical worker castes (Worker polymorphism) – many species have major (soldier), and minor castes with different specializations. E.g. soldier – defense, seed milling, abdominal food storage.
- Worker Age Castes – (Temporal Polymorphism) – workers of different ages perform different duties (e.g. young ants tend the brood, older workers forage (much more dangerous).
- Behavioral Castes – groups of similar individuals perform the same set of tasks within a given period.
- Division of labor is flexible and elastic – E.O. Wilson (1976) found when the fraction of minors is small, majors engage in the tasks usually performed by minors and efficiently replace them.

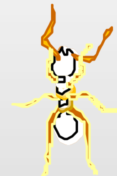

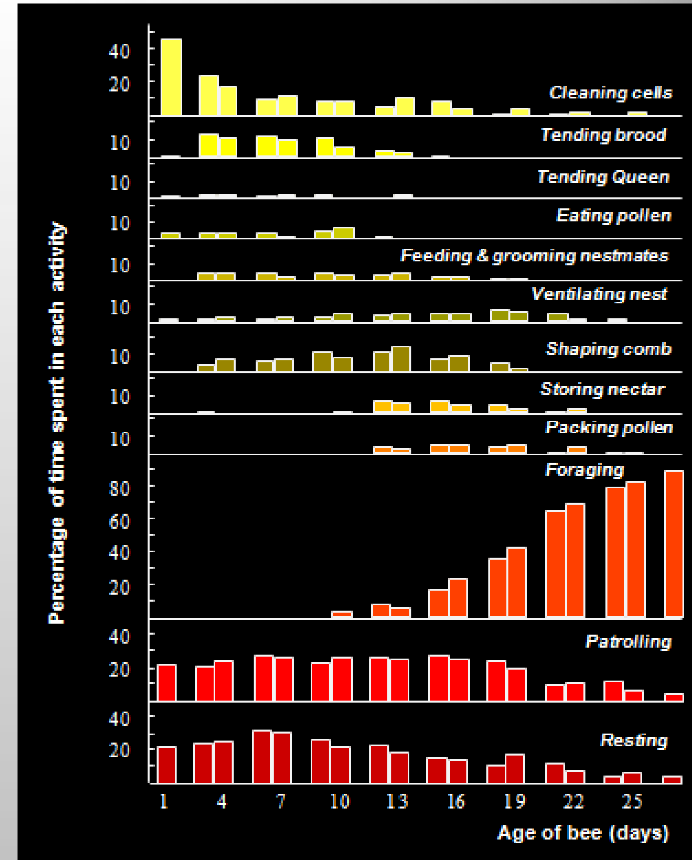

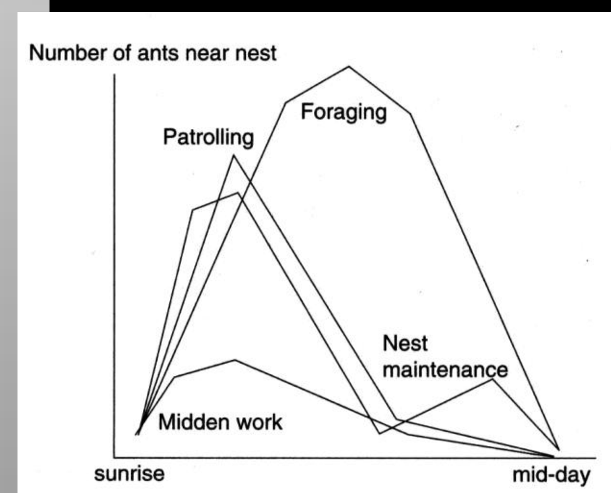

From D. Gordon, “Ants at Work”, 1999

Division of labor via Response Threshold Model

Individuals engage in a task when the task stimulus exceeds their threshold.

$$T_{\theta_i}(s) = \frac{s^2}{s^2 + \theta_i^2}$$

- minor ants are removed
- stimulus for minor tasks rises
- majors with a higher threshold start to perform minors tasks.

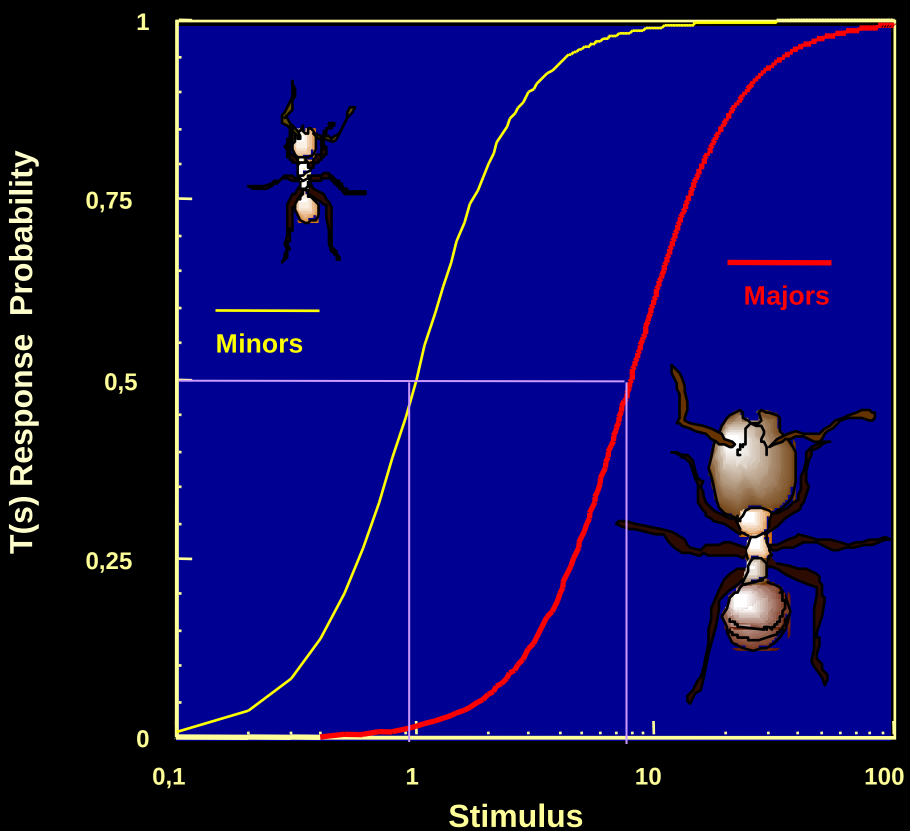

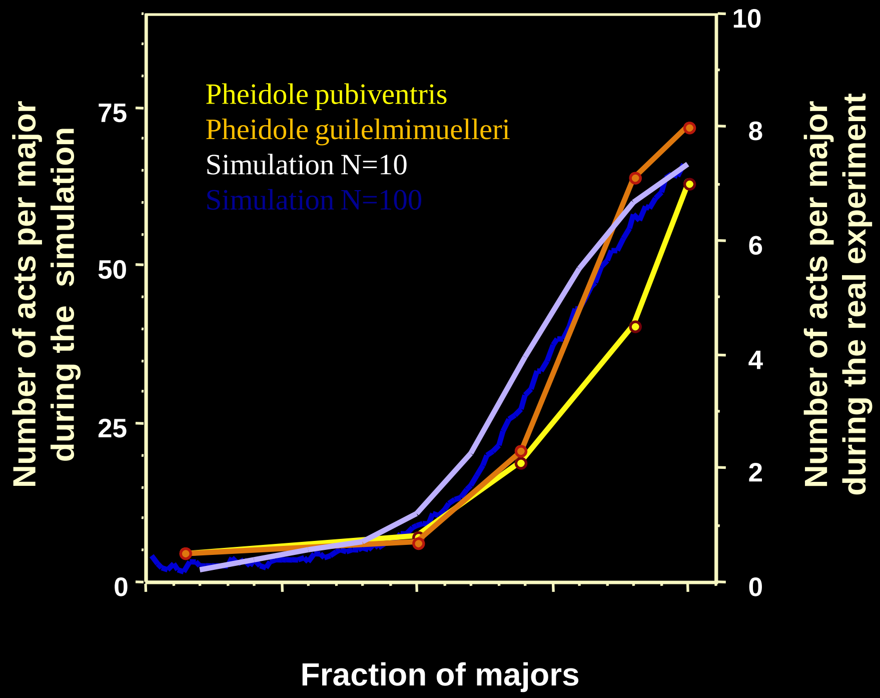

θ_{minors} θ_{majors}

s : intensity of the stimulus associated with the task

θ_i : response threshold

- Thresholds can vary with both age, and task experience, leading to "learning" and specialization.

Flocking and Collective Movement

- Occurs in all media (air, water, land), and many species (fish, birds, insects, mammals)
- Characteristics:
  - Rapid directed movement of the whole flock
  - No collisions
  - Reactivity to obstacles
  - Automatic coalescing and splitting of flocks
  - Tolerant of movement within flock
  - Tolerant of loss or gain of members
  - No dedicated leader
- Benefits:
  - Predator defense
  - Energy saving via turbulence reduction (fish, birds)

Artificial Flocking

- Early work by Craig Reynolds on "Boids"(1987) trying to implement natural looking computer animation of bird flocks.
- Basic Rules:
  - Avoid **obstacles**
  - Avoid **collisions** with nearby flockmates
  - Match **speed and direction** with nearby flockmates
  - Stay close to nearby flockmates (**flock centering**)
- Prioritize rules execute highest priority (conflict resolution).
- Sensing parameters determine "fluidity" of motion
  - Vision: inverse square, angle
  - Fish: pressure – inverse cube

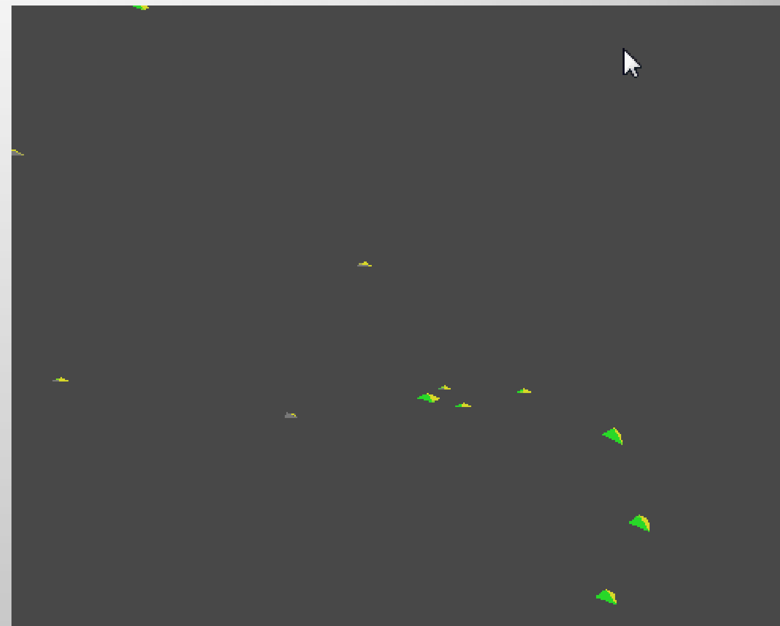

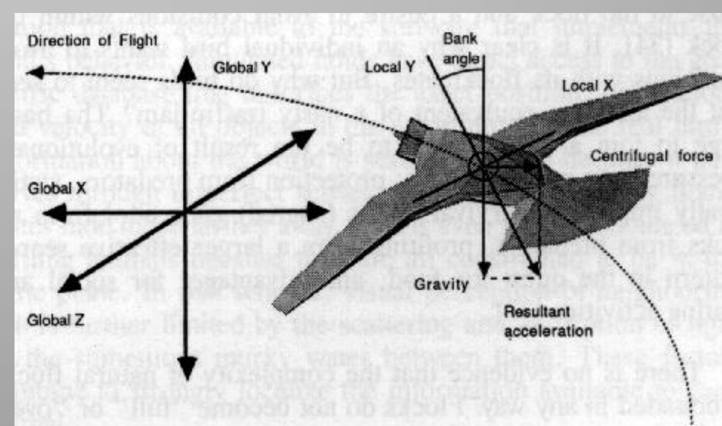

- Flight Model:
- Orientation
- Momentum
- Max Acceleration
- Gravity
- Lift

Swarms - Artificial

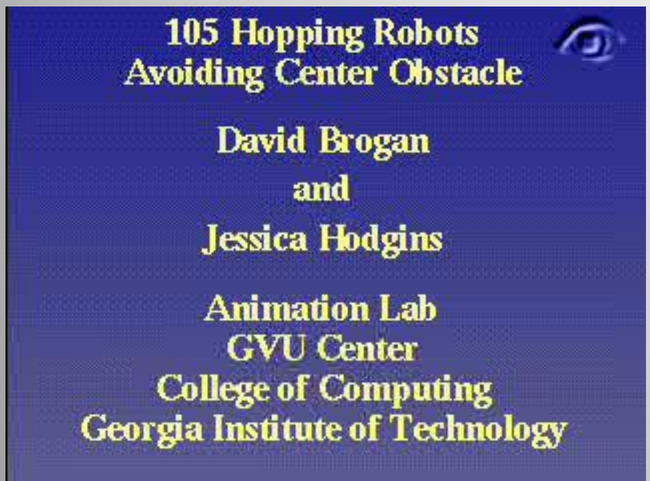

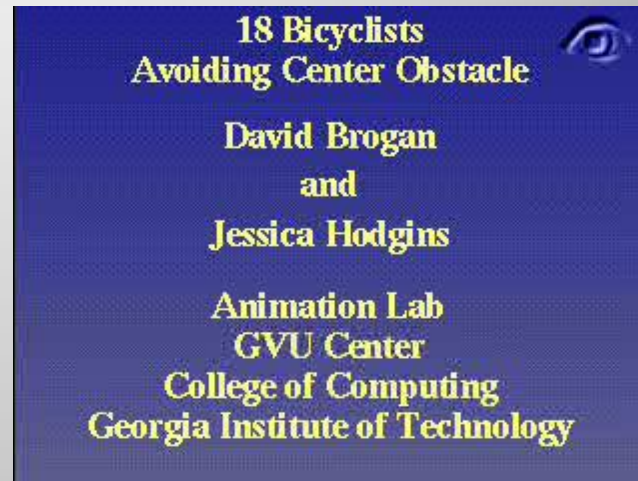

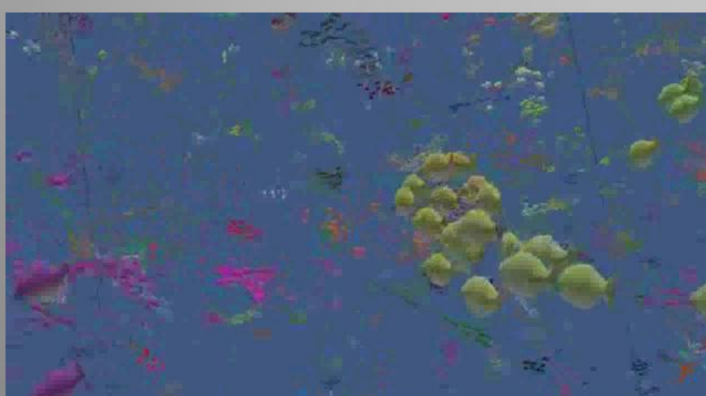

Craig Reynolds Fish on PS3

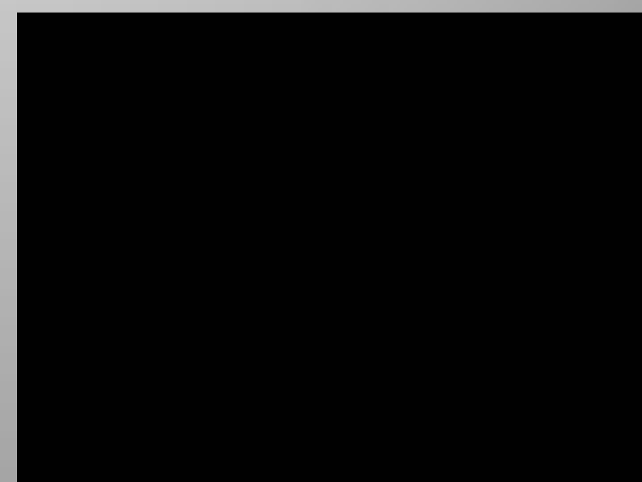

Steam Boids

The Movies!

Early Examples:

ANTZ

Lion King
“Wildebeast Stampede”

- These techniques are now routinely used in movie animations, and are incorporated into high end tools and Game Machines (PS3).
- MASSIVE is the current industry state-of-the-art.

Pedestrian and Traffic Planning

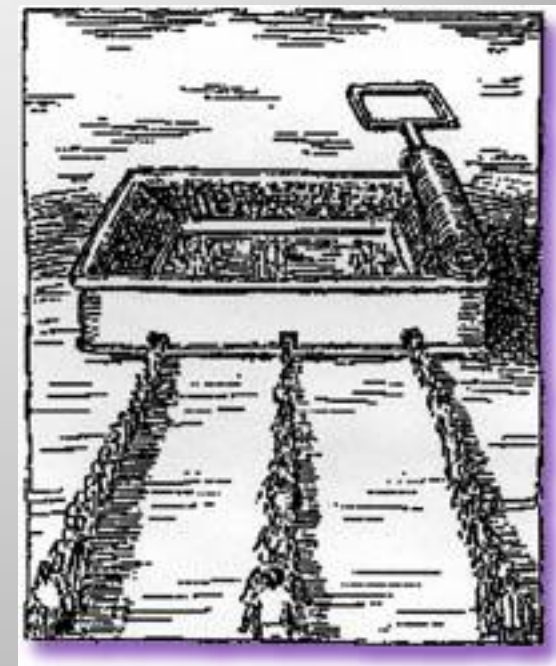

Dilemma Zone (Green/Yellow Transition) Detection and Warning System.

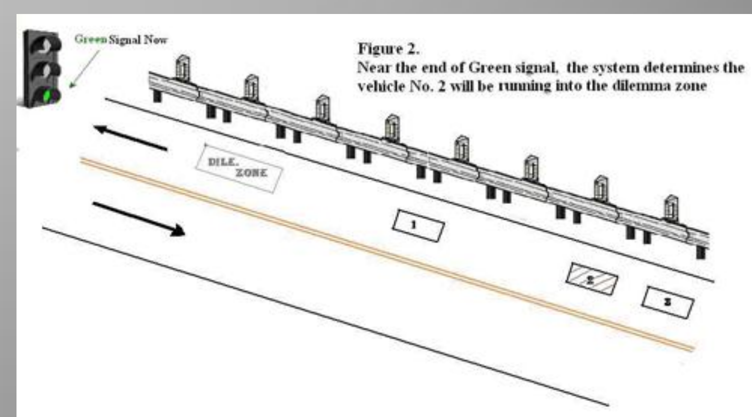

Engineering Swarm Intelligence -Robotics

Swarm Robotics

- Collision Avoidance Sensors
- Limited Local Communication (RF)
- Range and Bearing to neighbors critical for flocking etc.
- Simple Algorithms.
- Robust to unit failure.
- Autonomous operation of individuals.
- Learn the Controller
- Stigmergic communication through manipulation of the environment

- Approach:
    - Simulation
    - Probabilistic Modeling
    - Real Robot experiments

Puck Pushers: A Key First Achievement using Real Robots (Beckers, Holland, and Deneubourg, 1994)

Biological inspiration from clustering of dead ant corpses.

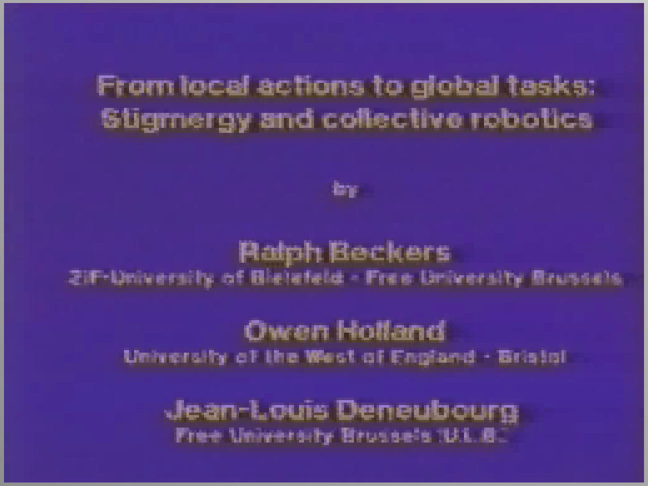

Probabilistic Modeling via cluster modifying probabilities:

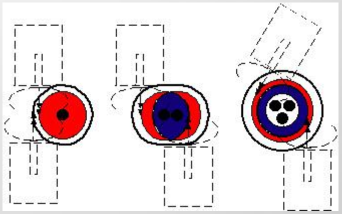

Martinoli A., Ijspeert A., and Mondada F. , *Understanding ollective Agregation Mechanisms: From Probabilistic Modeling to Experiments with Real Robots*. **Robotics and Autonomous Systems**, 29:51-63,1999

Kazadi S., Abdul-Khaliq A., Goodman R., *On the Convergence of Puck Clustering Systems*. **Robotics and Autonomous Systems**, 38 (2), pp. 93-117, 2002.

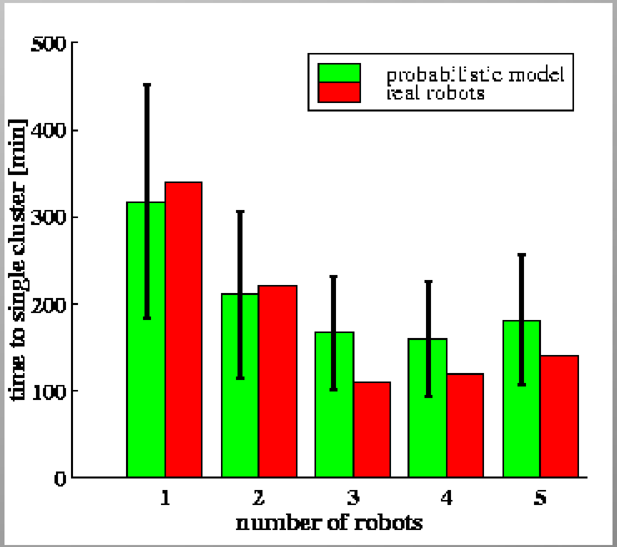

Khepera Stick-Pulling

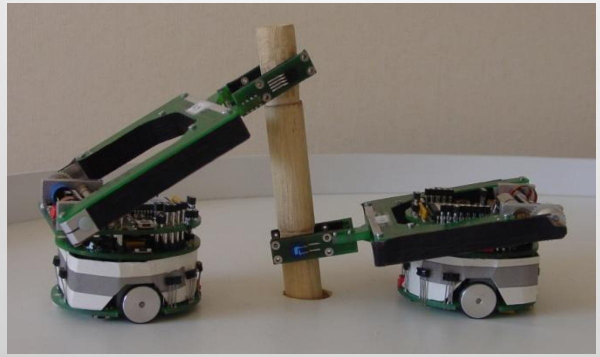

Collaboration via Indirect Communication

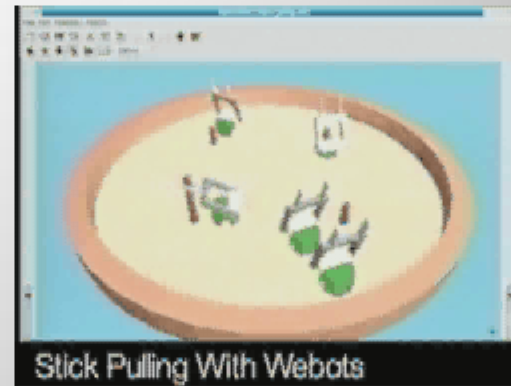

Realistic Simulation

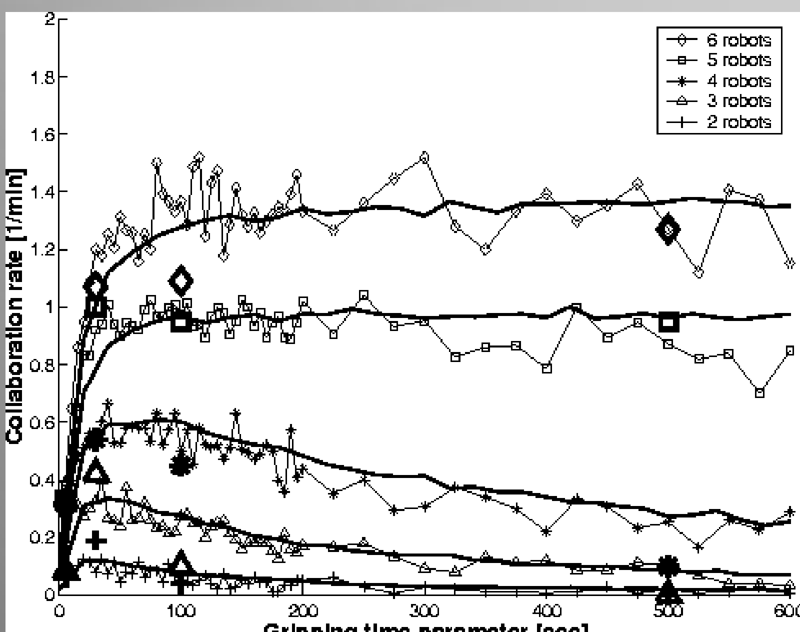

Real Robots

- [Martinoli and Mondada, ISER, 1995]
- [Ijspeert et al., *AR*, 2001]

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

Khepera Aggregation Experiments

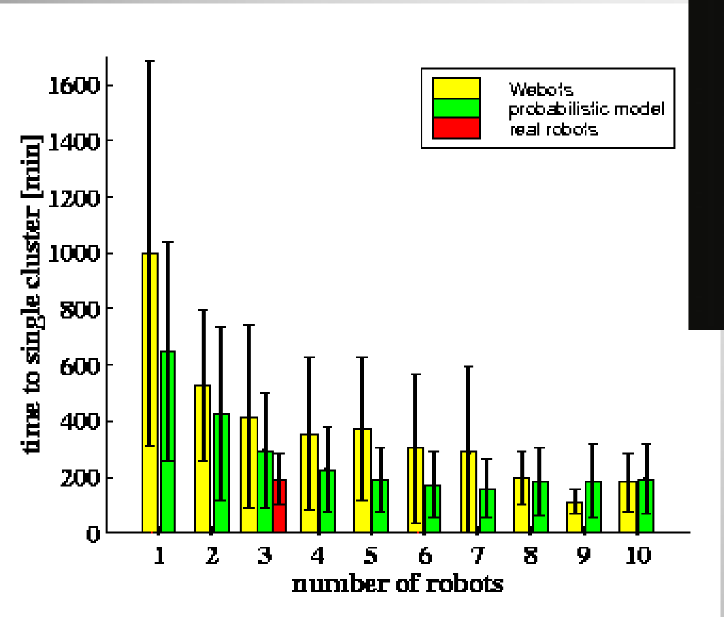

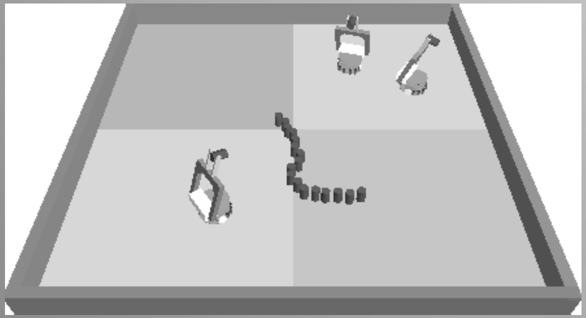

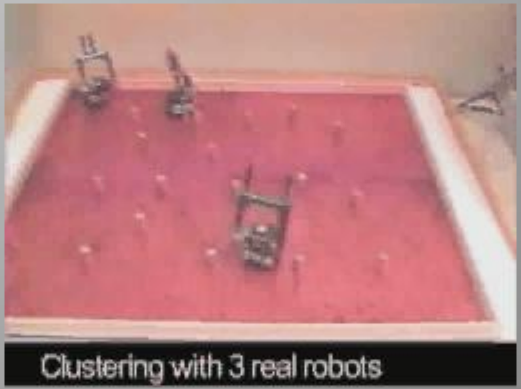

IDSIA

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

Multi-Level Modeling Methodology

$$\frac{dN_n(t)}{dt} = \sum_{n'} W(n \mid n', t) N'_n(t) - \sum_{n'} W(n' \mid n, t) N_n(t)$$

Increasing Abstraction

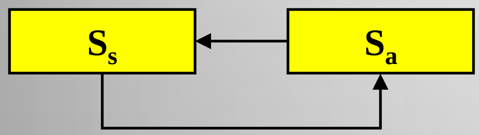

Macroscopic: rate equations, mean field approach, whole swarm

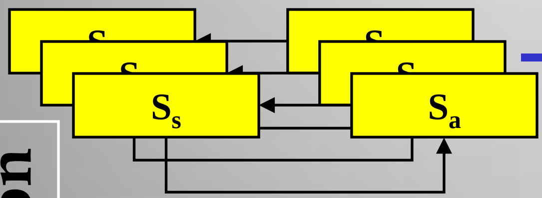

Microscopic – Agent-based: multi-agent models, only relevant robot features captured, 1 agent = 1 robot

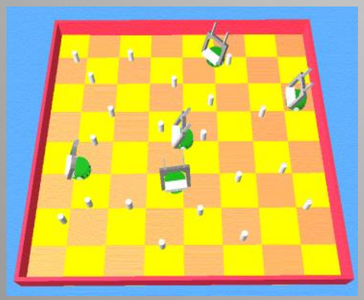

Microscopic – Module-based: intra-robot (e.g., S&A, transceiver) and environment (e.g., physics) details reproduced faithfully

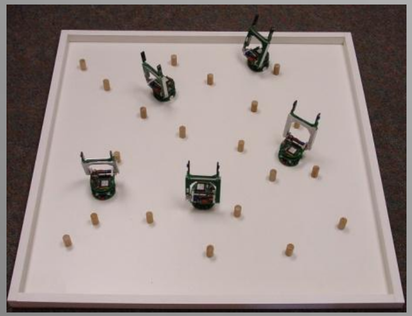

Target system (physical reality): info on controller, S&A, communication, morphology and environmental features

1. Probabilistic FSM description for environment and multi-agent system; arbitrary state granularity
2. Semi-Markovian properties: the system future state is a function of the current state (and possibly of the amount of time spent in it)
3. Nonspatial metrics for swarm performance

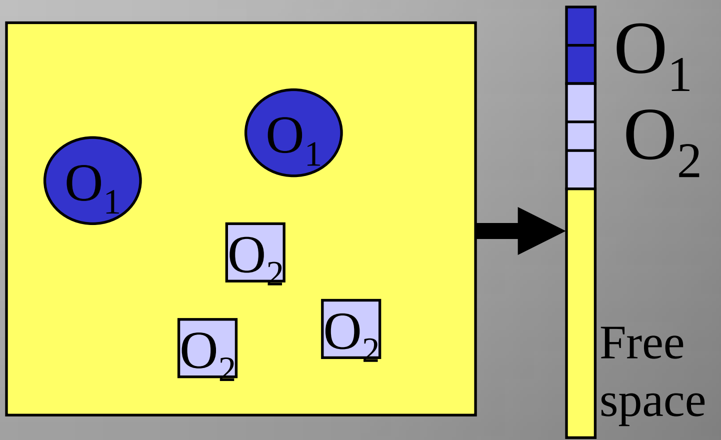

2D physical space -> 1D prob. space

Robo Sheepdog

The first animal-robot interaction (1998)

Moorebot Flocking (Caltech)

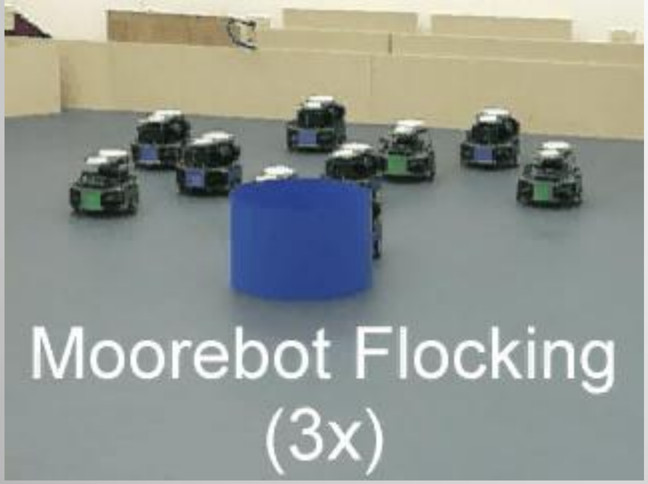

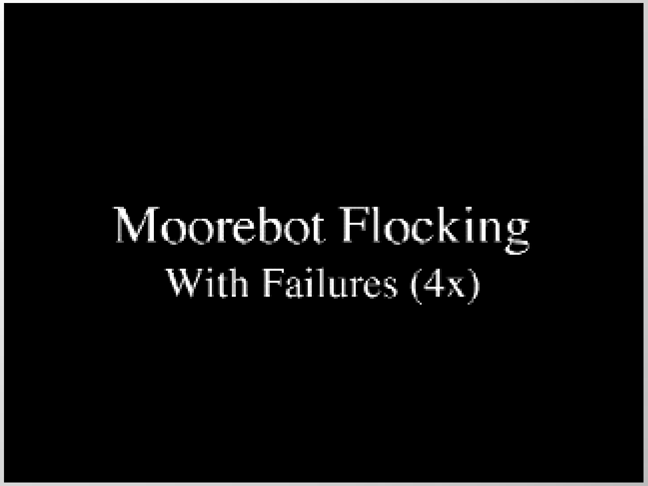

Pseudo Flocking in Traffic

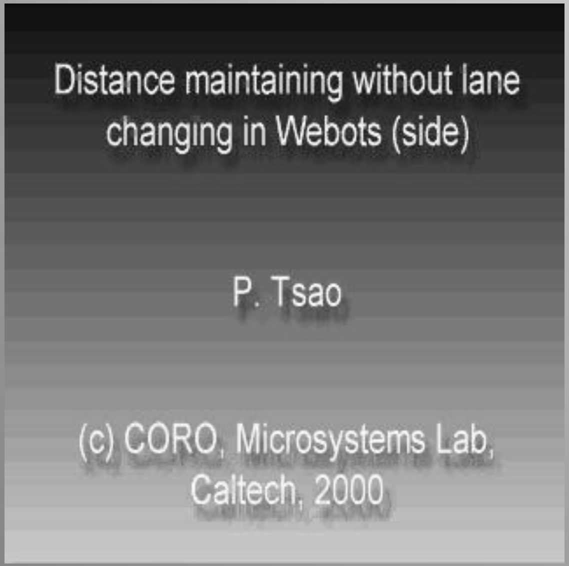

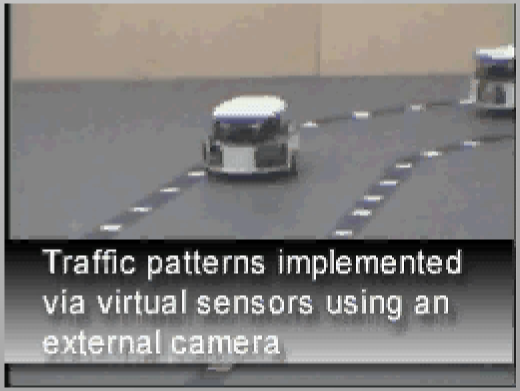

The Flying Flock (UWE)

- Owen Holland
- Alan Winfield et. al.
- Bristol Robotics Lab
- University of the West of England-University of Bristol

Swarm System Ltd wins UK Grand Challenge Award for Most Innovative Idea on 15-16 August 2008!

Technology

The Swarm Systems technology is a swarm of OWL quadrotor MAVs and a Ground Station. The technology elements include:

Air Vehicle

quadrotor
around 1kg in weight
standard and carrier wave GPS
sonar
inertial navigation unit
magnetometer
altitude pressure sensor
powerful onboard processor running Linux or Windows

Sensors

visible 12 megapixel camera
video

Ground station planner
mission controller
target recognition software
report generator

UK Defence Procurement Minister Lord Drayson congratulating Swarm Systems CEO Stephen Crampton (left) and Prof. Owen Holland (right) on winning MOD funding for the Grand Challenge

www.swarmsys.com

Odor Source Localization (Caltech)

- Given odor plume, find the source of the odor plume as efficiently as possible.
- Chemical Agent Tracking
- Task Decomposition
  - Plume finding
  - Plume traversal
  - Source declaration

MooreBot with Integrated Wind and Odor Sensors

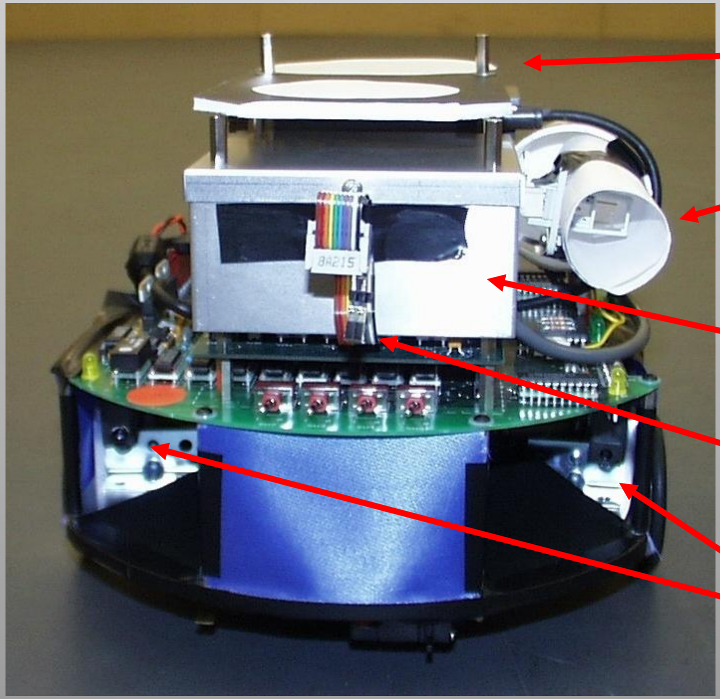

Collective Plume Tracing

- Single Robot Behavior
  - Spiral-Surge Algorithm
  - Loosely based on moth “casting”
    - If no hit – spiral out
    - If hit – surge upwind

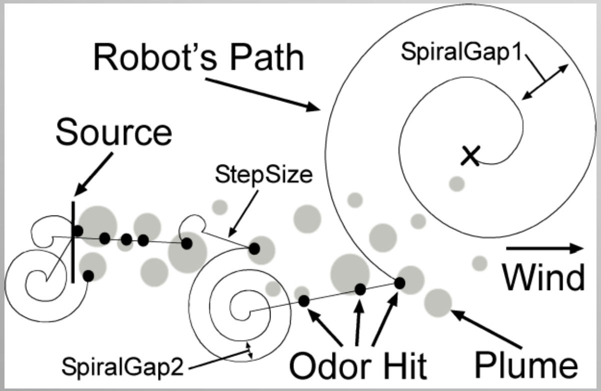

- Multi-Robot Collaboration
  - If no hit – *attracted* to nearest robot with hit.
  - If no other hits spread out (*repulsed*).

Optimization via Reinforcement Learning

- Define performance metric = Group Energy + Time First Robot.
- Optimize parameters (spiralgap size, surge length, etc) relative to metric using off-line realistic simulator.
- Learned solution (p3) significantly better than hand-coded ones (p1,p2)

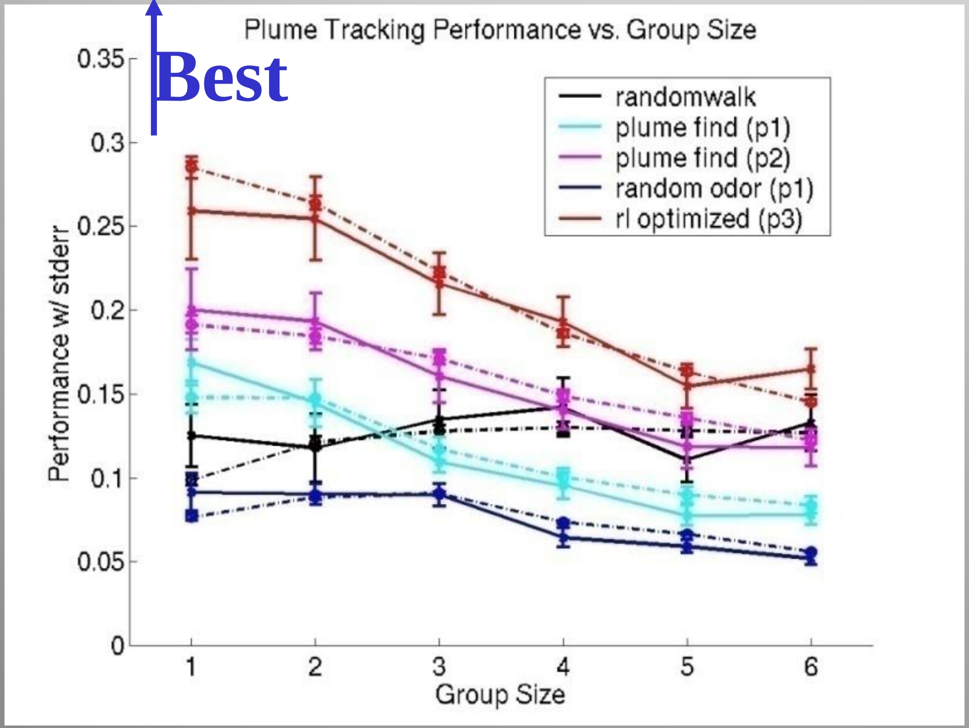

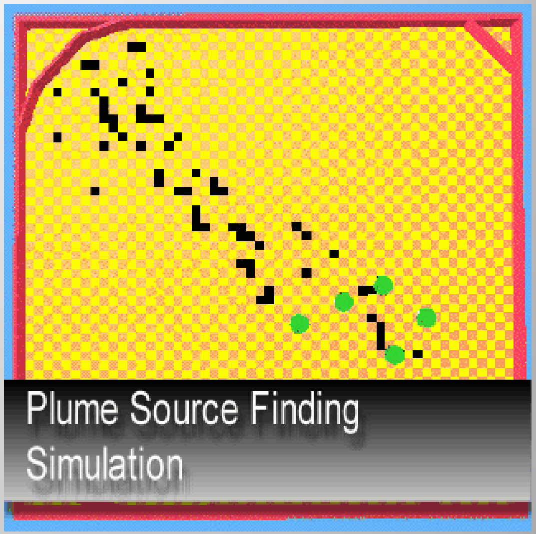

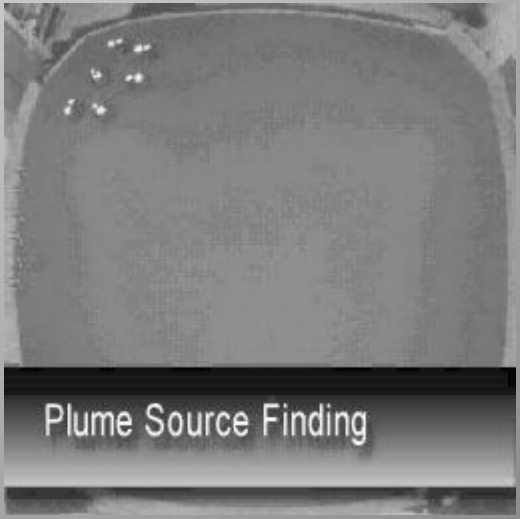

6 Robots

Collaborative Radiation Interdiction (Caltech) and Smiths Detection (Pasadena)

- Handheld CZT Rad Detector
- Displays Gamma ray Spectrum, Direction to Source
- Collaborative search to localize source via real-time networking
- Multi-player game simulator: Half-Life2
- Bayesian inference probabilistic multi-scale “heat map” in real time

Matt Wu, et al
InfoSpheres Research Group   *www.infospheres.caltech.edu*
Caltech

- **Where we are at in SI:** Using the biology - theory - experiment - application path SI has evolved over two decades into a into a reasonably successful multidisciplinary science and engineering effort.
- **Results:**
  - Optimization Algorithms, Communications .
  - Real Robotic Systems just becoming technically feasible.
- **Lessons:**
  - Real robot experiments are very very time consuming ! And expensive!
  - Probabilistic modeling is accurate but very time consuming on intellect.
  - Latest physics based simulators from the games community (Nvidia PhysX, Half-Life2, Player-Stage, are now so good that simulation is the way to go.
- **Challenges:**
  - Understanding the biology at a finer granularity, to understand the basis of behaviors.
  - Energy Autonomy.
  - Many Swarm robotic applications in unstructured land environments, and in the sea.
- **On-going work:**
  - Develop robot controllers that run a internal model (simulation) of the environment, and an internal model of the "self", to predict and execute action. A consciousness?

Contact me at : rod@goodman.name or www.rodgoodman.ws