ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Distributed Intelligent Systems

EPFL, WS 2017-2018

http://disal.epfl.ch/teaching/distributed_intelligent_systems/

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Distributed Intelligent Systems – W1

Part I: Course Organization, Team, and Content

Team beyond this Course

http://disal.epfl.ch

- **Instructor:** Alcherio Martinoli
- **Guest lecturer**: Ali Marjovi

- **Teaching assistants:**
  - Faezeh Rahbar (Head TA, PhD student)
  - Alexander Bahr (TA, postdoc)
  - Ali Marjovi (TA, postdoc)
  - Anwar Quraishi (TA, PhD student)

- **Support staff:**
  - Lukas Pestalozzi (help TA, SIN master student)
  - Bahar Haghighat (PhD student)
  - Zeynab Talebpour (PhD student)

Access to e-material, exercise, and building room

- Moodle web site for the course:
  - Students registered for the course on IS-Academia are automatically registered for the course on moodle (re-synch daily over night)
  - If issues in accessing the moodle web site despite registration in ISA, please contact Head TA Faezeh Rahbar
- For non-SIE students we need to request for you explicit access rights for the computer room and building; we will do so based on the final enrollment list (frozen on September 29)

Rationale for This Course

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Rationale

- Well-balanced course: theory, algorithms and experimental labs
- Understand quantitatively natural collective phenomena (focus on biological societies) and how to combine bio-inspired principles with advanced engineering methods
- Understand how to model, design, control, evaluate, and optimize distributed intelligent systems
- Learn to process scientific literature efficiently: prioritize readings, dig out papers, find connections

What is this Course about

- Distributed natural and artificial systems
- Coordination algorithms
- Distributed sensing and action
- Models, simulation tools, and machine-learning targeted to distributed intelligent systems
- Multi-robot systems and wireless sensor networks

Course Prerequisites

- C and Matlab knowledge
- Fundamentals of programming
- Fundamentals of probability calculus
- Fundamentals of analysis (differential equations, continuous and discrete time)
- Fundamentals of linear algebra
- Fundamentals in signal and systems

For SIE students: BS introductive course on Signals, Instruments, and Systems highly recommended!

Organization of the Course

This Edition

- Preserved good innovations from last edition (AY 2015-2016): lab verification tests, primary/secondary reading, limited number of well-prepared TAs in the lab sessions
- Major difference: no course project
- Further goal: labs better aligned with lab verification tests and a bit more oriented towards concrete implementation
- dis-ta@groupes.epfl.ch for any issue (e.g., inquiries, office hours, etc.)

Credits and Workload

- 5 ECTS
- 1 ECTS = 30 h workload → 150 h workload
- Rough breakdown
  – 75 h lecture (including reading and exam prep)
  – 75 h exercise (including preparation + tests)

Grade

- Final written exam, winter session:
  - 180 minutes;
  - open book with simple non-programmable calculator;
  - all topics covered in the lecture/exercise and selected distributed reading material
- 50% performance during semester, 50% performance during the exam (compromise US/Europe style)
- During semester: weighted average of two lab verification tests; first test: 30% (verifies content acquisition of 6 labs); second test: 20% (verifies content acquisition of 4 labs)

Lecture

- Tue 10:15-12:00
- This week and last week exceptionally also on Wednesday (13:15 – 16:00) instead of exercises in MA B1 11 (84 seats)

Lecture Notes

- Preliminary lecture slides in pdf format available for download on Moodle before each lecture (Monday late evening), definitive ones after lecture by Friday at latest)
- Will notify when ready in definitive format via e-mail

Reading and Handouts

- Policy: master, research-oriented course → no manuscript! → slides + papers + web
- Break down in 3 categories:
  - Primary: covered substantially during the lecture; available on moodle
  - Secondary: covered marginally during the lecture; available on moodle
  - Tertiary: pointers on the lecture notes for interested students, not covered in the lecture and not available for download
- Roughly 45 single-column pages/week of primary literature to read; list and primary/secondary breakdown subject to change during the semester
- Primary and secondary reading distributed the week before for easing exercise preparation & lecture understanding

Suggestions for a Successful Course Material Processing

From last years experience:

- For high-gear courses such as this one with a lot of raw material to process: it is worth taking advantage of the lecture for having an idea about what’s important and what not
- Trained ability: reading what’s needed and quickly, seeing connections between various “raw” pieces of the puzzle

Labs

- Lab session: 3 h on Wed, 13:15-16:00, GR B0 01 and GR C0 02
- Mini-tutorial (< 10 min) by the main lab designer at the beginning of the lab
- 3 TAs per lab session (1 designer, 2 testers)
- 10 lab sets total, not graded (solution distributed)
- 2 lab verification tests, in the computer room, graded (personalized feedback), mixture of computer-based and paper-based exercises:
  - W8: lab 1-6
  - W13: lab 7-10

Suggestions for a Successful Exercise Series

From last years experience:

- Read the lab assignments in advance, in this way you will be more efficient when the TAs are around for helping you on the toughest questions …
- Have an idea of the point distribution of any assignment: this roughly correspond to the breakdown in time you should have; if your time is tight invest where it is worth!
- Take lab notes so that you will find them for the lab verification tests
- If you do not work enough independently during labs, it will be difficult to solve problem set alone in the tests
- “Paper-based” questions a good training for the final exam
- Previous edition assignment are on the web (with all the remaining material)

Collaboration Policy

- Lecture and exam preparation: encouraged
- Lab: discussion encouraged but work individually
- Lab verification tests and final exam: collaboration penalized …

Course Syllabus and Summary

Goal

- Course overview
- Course flavor
- 5 main blocks
- A few slides per block

Block I – Swarm Intelligence

- Key Principles of Swarm Intelligence
- Trail laying/following mechanisms
- Ant Colony Optimization as an example of a successful multi-agent metaheuristic

From Natural to Artificial Systems

- Modeling to understand microscopic to macroscopic transformation
- Modeling as interface to artificial systems

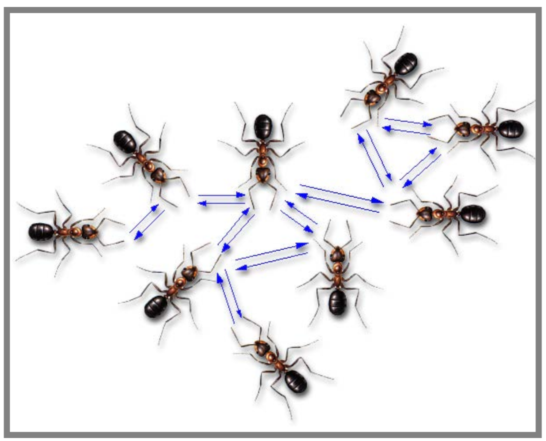

Individual behaviors and local interactions

Modeling

Ideas for artificial systems

Global structures and collective decisions

Biological Mechanisms and Models

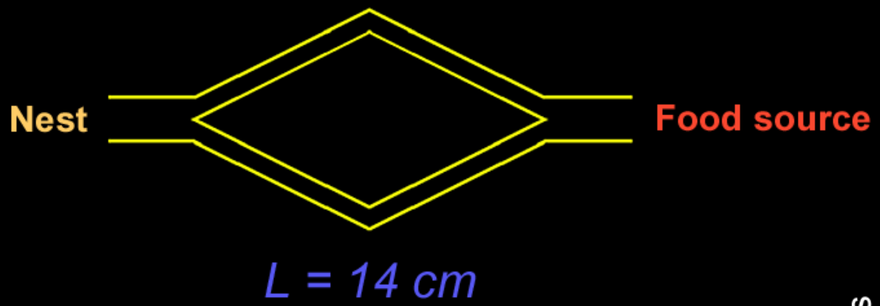

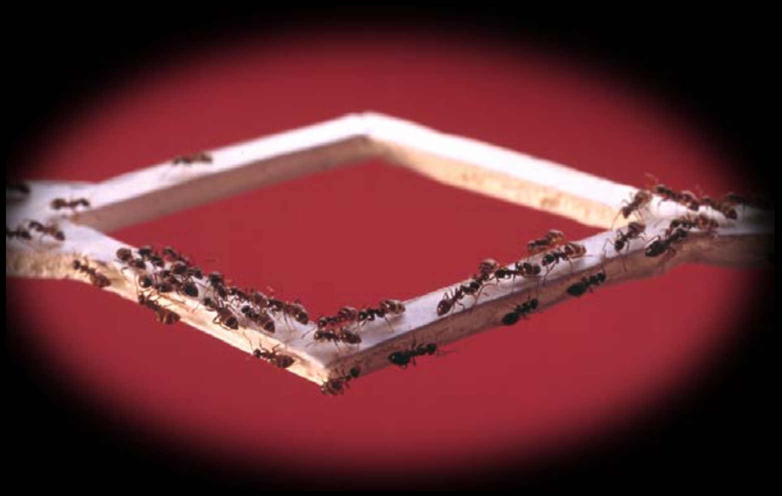

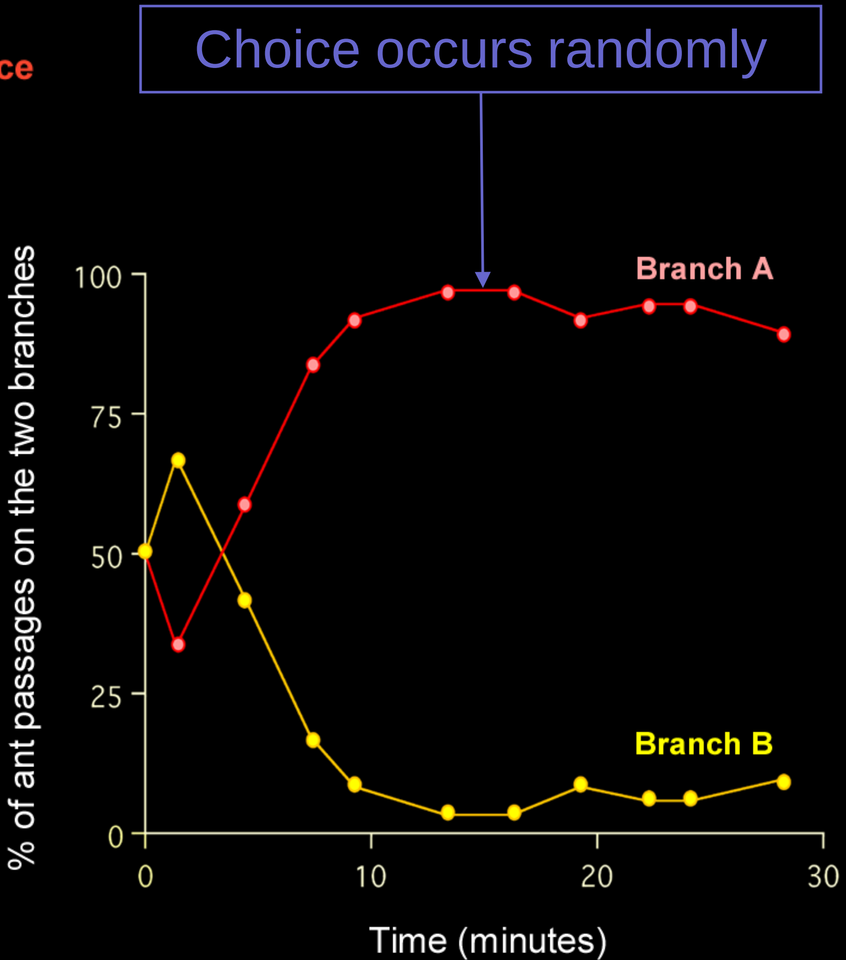

(Deneubourg *et al.*, 1990)

The Traveling Salesman Problem

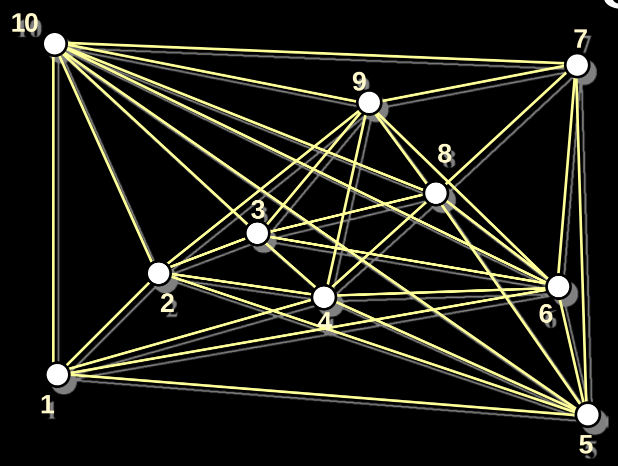

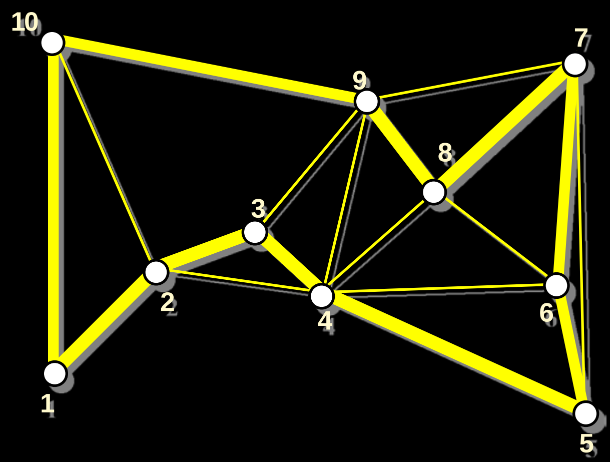

Graph (*N*,*E*)
N : set of cities (nodes)
E : set of connecting roads (links)
d_{ij} : distance between city i and j

Problem: Find the shortest path which allow the salesman to visit once and only once each city in the graph

Difficulty: NP-hard problem; time for computing the shortest route grows in a nonpolinomial way with the number of cities in the network -> metaheuristics provide near-optimal solutions!

Block II – Individual Nodes

- Introduction to mobile robotics
- Robotic tools (simulation and real HW)
- Basic control architectures

Real and Simulated e-puck

- Appropriate size for desktop
- Multi-robot operation ok
- No manipulation, no highly accurate odometry

- Webots realistic robotic simulator
- Discrete sensor and actuators
- Single and multi-robot simulator

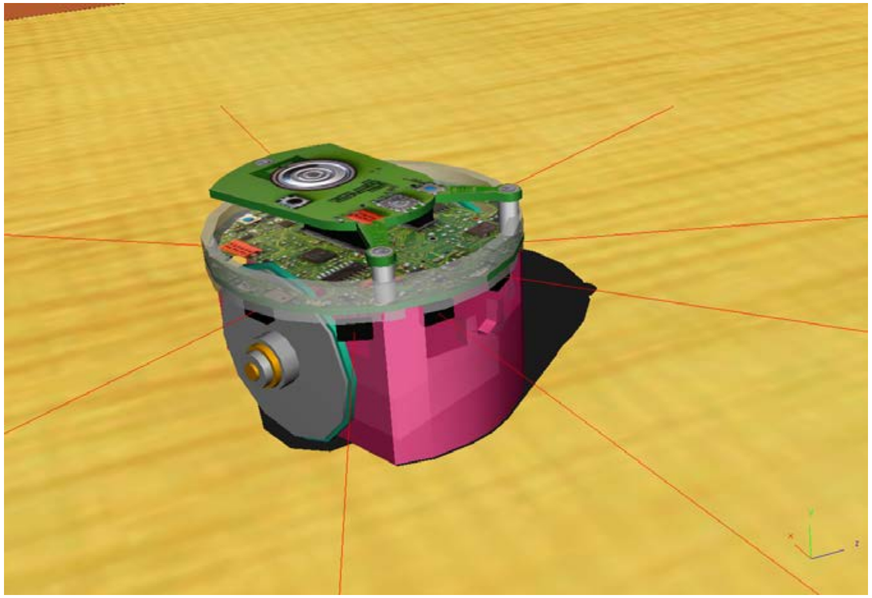

Perception-to-Action Loop

- Reactive (e.g., nonlinear transform, single loop)
- Reactive + memory (e.g. filter, state variable, multi-loops)
- Deliberative (e.g. planning, multi-loops)

• sensors

• actuators

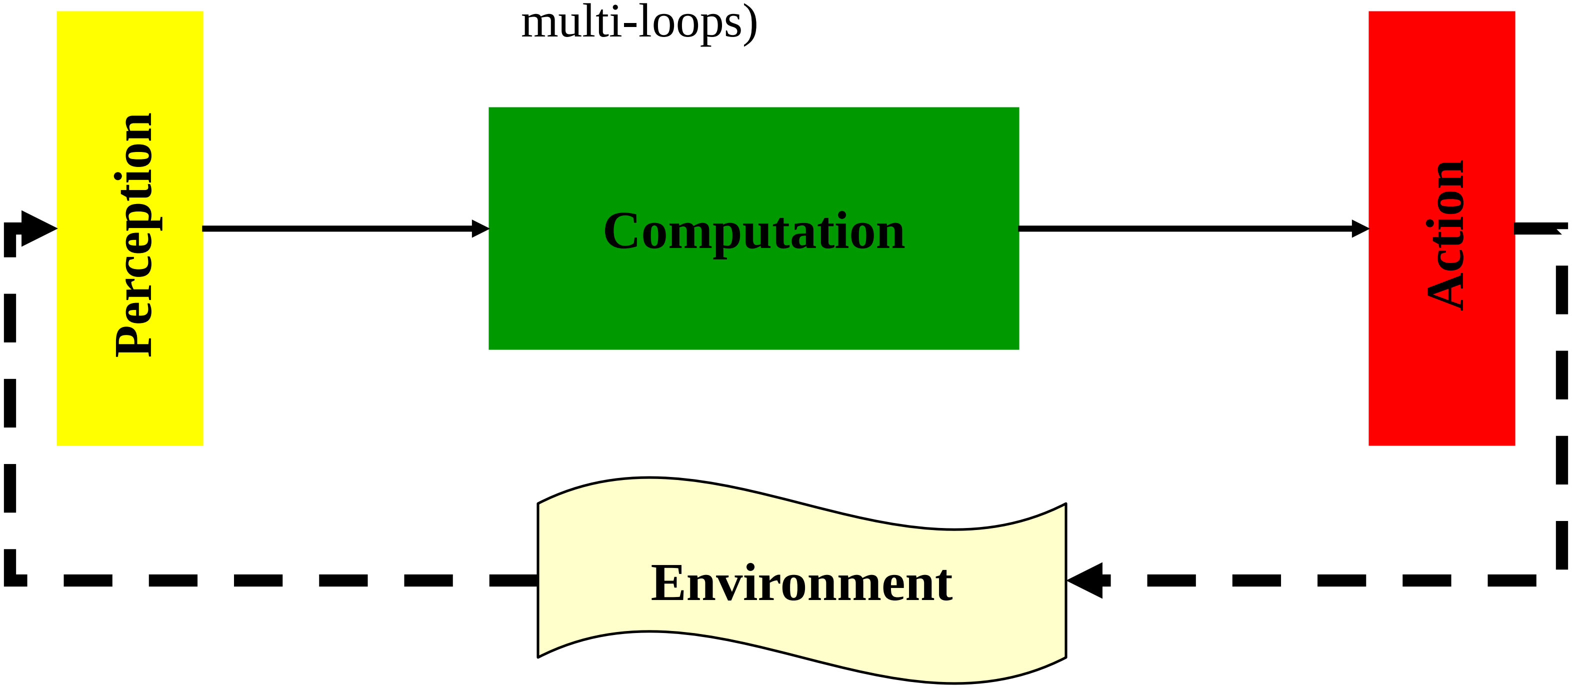

Robot Localization

- Key task for:
  - Path planning
  - Mapping
  - Referencing
  - Coordination

?

N 46° 31’ 13”
E 6 ° 34’ 04”

- Type of localization
  - Absolute coordinates
  - Local coordinates
  - Topological information

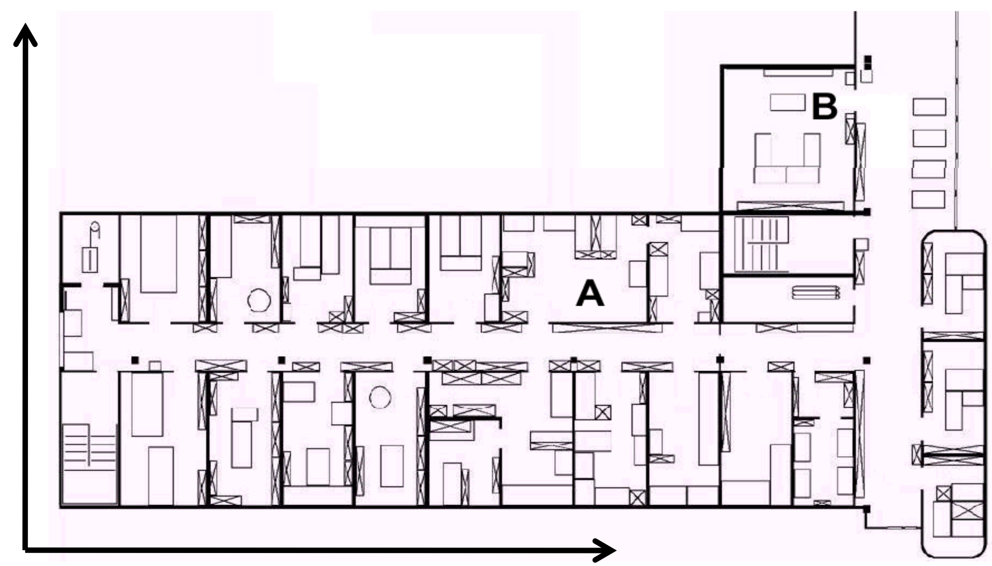

Block III – Coordination Algorithms

- Collective movements and spatial consensus
- Division of labor and task allocation
- Collective decisions

Ex. of Collective Movements

Reynolds'Rules for Flocking

1. Separation: avoid collisions with nearby flockmates

2. Alignment: attempt to match velocity (speed and direction) with nearby flockmates

3. Cohesion: attempt to stay close to nearby flockmates

Position control

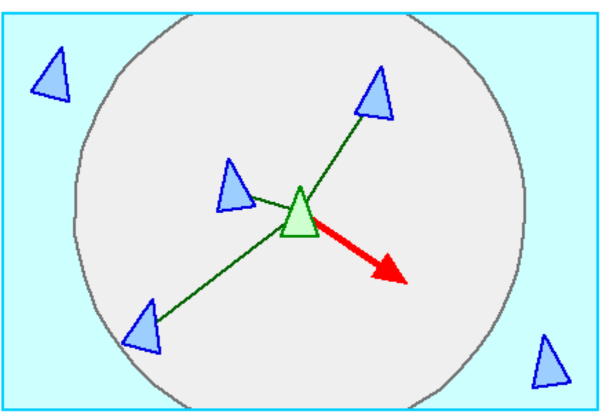

separation

Velocity control

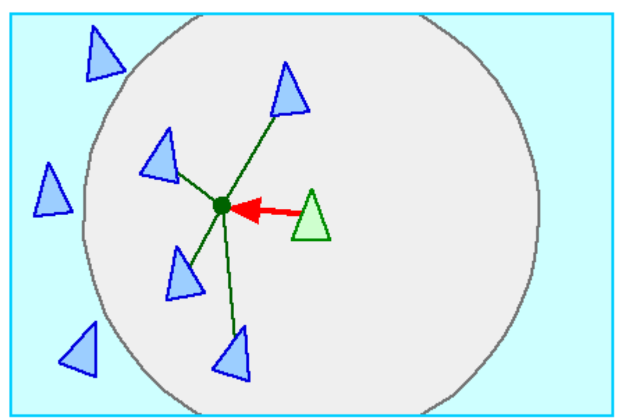

alignment

Position control

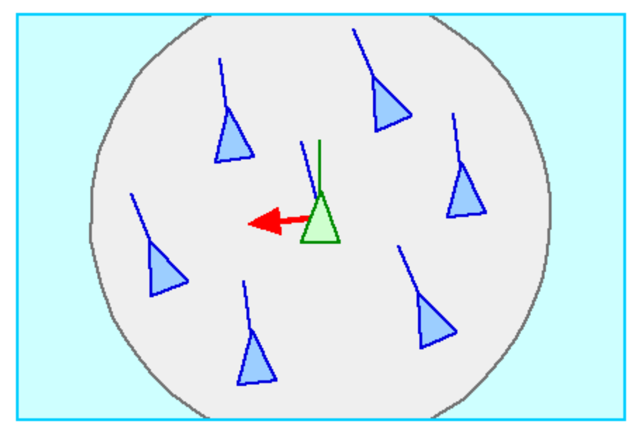

cohesion

Implementation of Flocking Rules in Artificial Embedded Agents

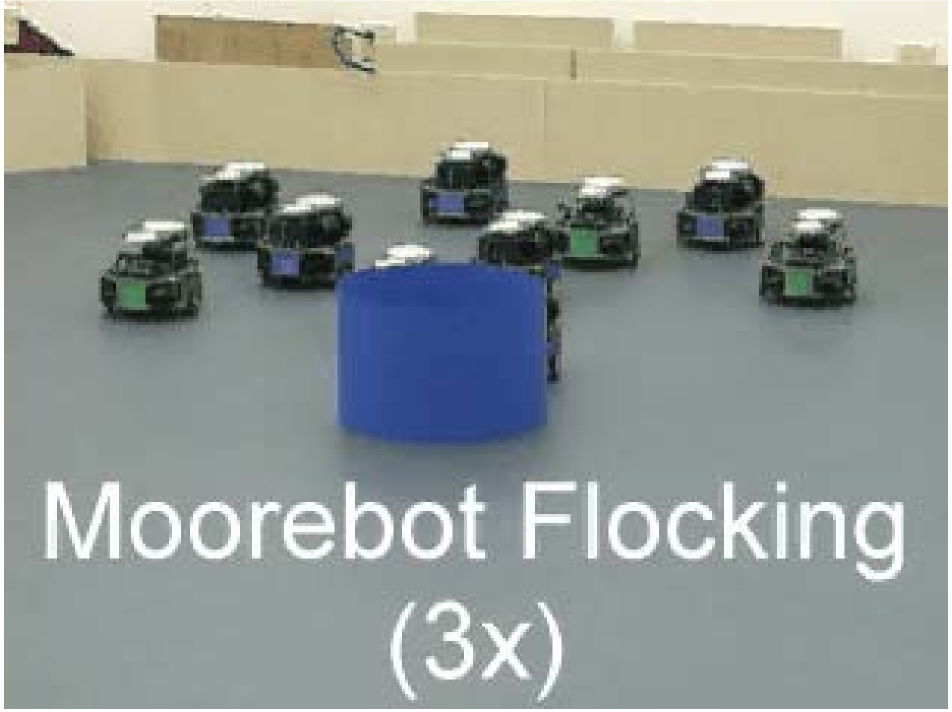

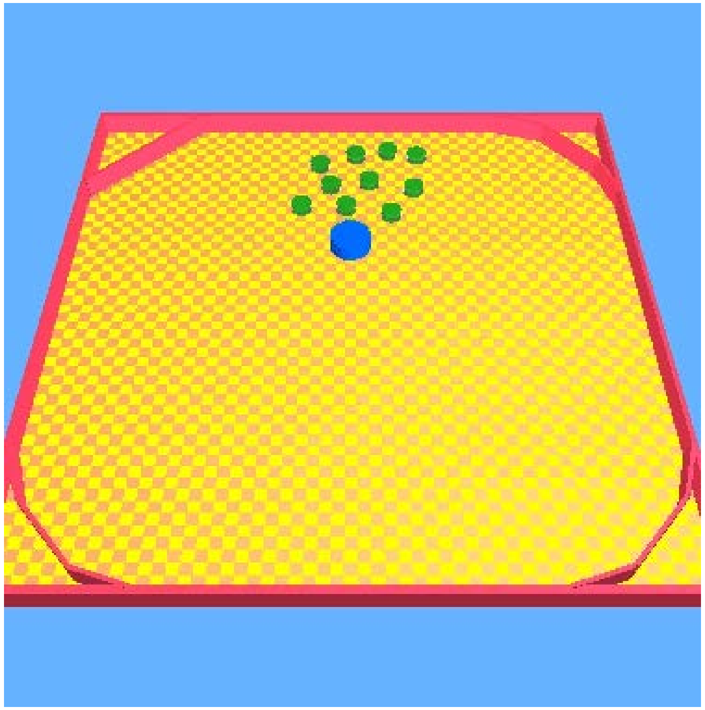

Real robots

Realistic simulator (Webots)

Market-Based Coordination

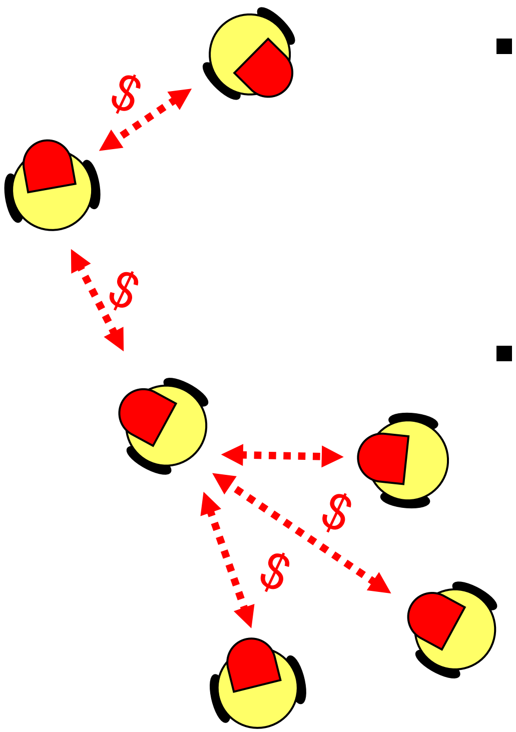

Robots simulate a market economy:

- Tasks, resources are commodities of measurable worth.

- When robot performs task:

  gets paid for service it provided (+ $)

  pays for resources it consumed (- $)

- Robots trade tasks and resources to maximize profit

Idea! pursuit of individual profit leads to efficient team solutions.

✓ Robust, fast, handle complex tasks

✓ Can take advantage of centralized planning

MICAz

- Atmel ATmega128L
  - 8 bit microprocessor, ~8MHz
  - 128kB program memory, 4kB SRAM
  - 512kB external flash (data logger)
- Chipcon CC2420
  - 802.15.4 (Zigbee)
- 2 AA batteries
  - about 5 days active (15-20 mA)
  - about 20 years sleeping (15-20 µA)
- TinyOS

Set-up and Collective Decision Algorithm

- 10 robots execute wall-following behavior (CW or CCW, initially random)
- announce their current direction on the radio channel

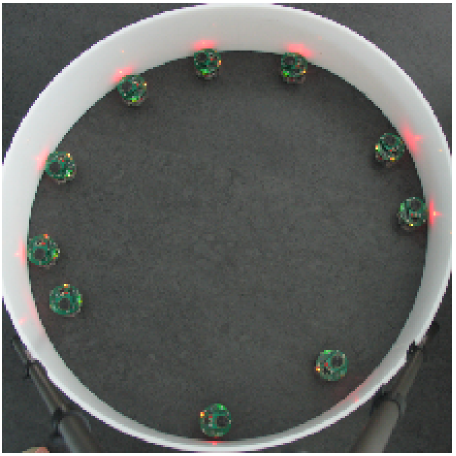

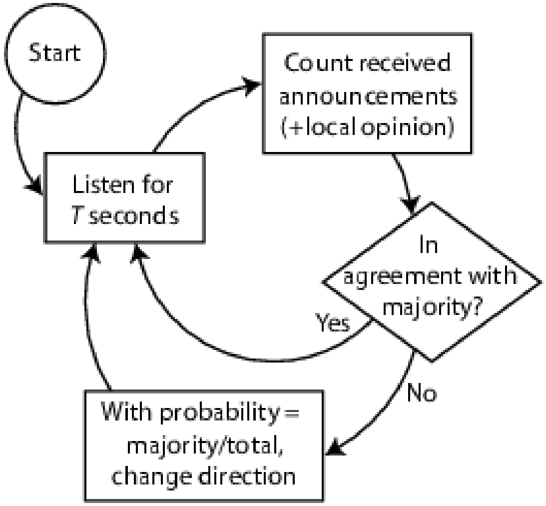

- # of votes not constant
- probabilistic decision
- communication range affects time to convergence.

[Cianci et al, SAB-SRW 06]

Selecting a Shelter

- ***Leurre*: European project focusing on mixed insect-robot societies (http://leurre.ulb.ac.be)**

- A simple decision-making scenario: 1 arena, 2 shelters
- Shelters of the same and different darkness
- Groups of pure cockroaches (16), mixed robot+cockroaches (12+4)
- Infiltration using chemical camouflage and statistical behavioral model

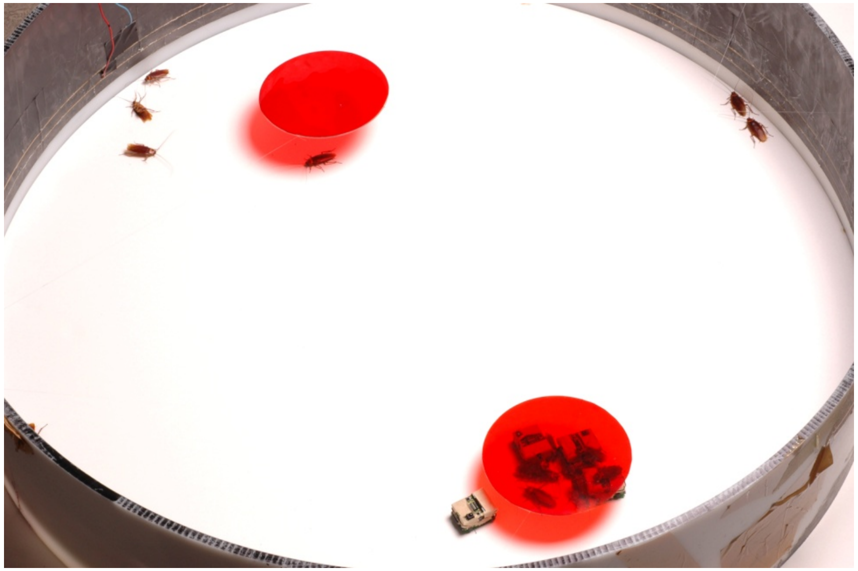

[Halloy et al., *Science*, Nov. 2007]

Block IV – Modeling and Optimization Methods

- Multi-level modeling
- Evaluative learning algorithms
- Particle Swarm Optimization as an example of a successful population-based learning algorithm
- Noise-resistance and distributed implementation

Multi-Level Modeling Methodology

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

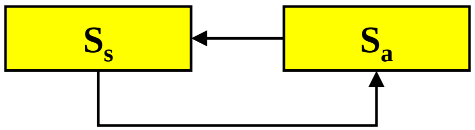

Macroscopic: rate equations, mean field approach, whole swarm

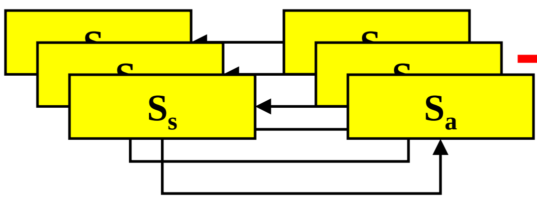

Microscopic : multi-agent models, only relevant robot feature captured, 1 agent = 1 robot

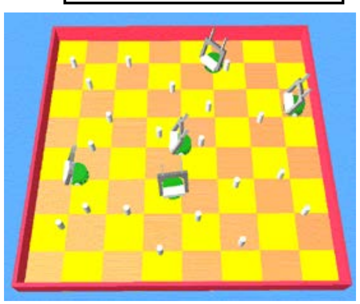

Submicroscopic : intra-robot (e.g., S&A, transceiver) and environment (e.g., physics) details reproduced faithfully

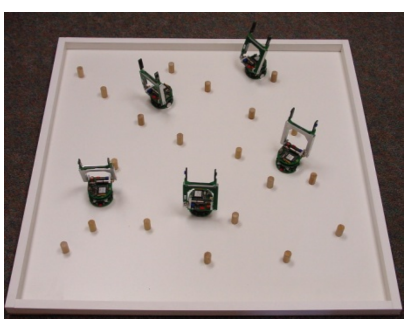

Target system (physical reality): info on controller, S&A, comunication, morphology and environmental features

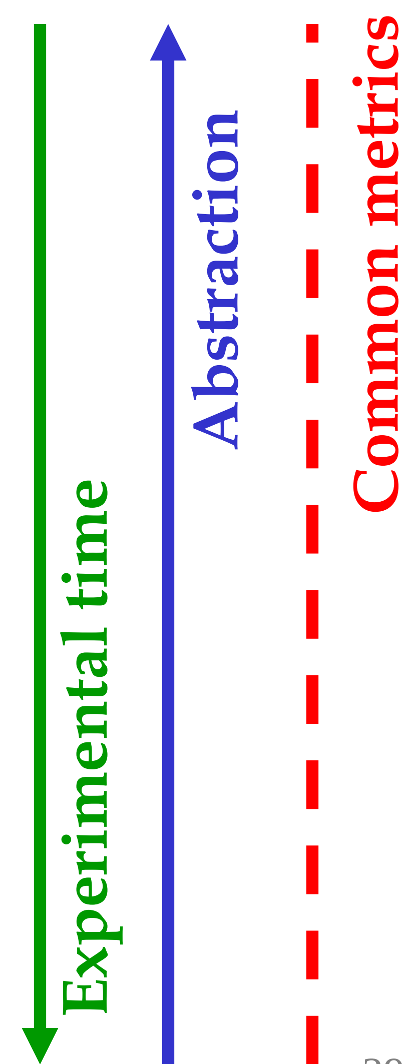

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

The Main PSO Loop
(Eberhart, Kennedy, and Shi, 1995, 1998)

At each time step *t*

for each particle i

for each component *j*

update the velocity

$$v_{ij}(t+1) = wv_{ij}(t) + c_p rand()(x^*_{ij} - x_{ij}) + c_n rand()(x^*_{i'j} - x_{ij})$$

then move

$$x_{ij}(t+1) = x_{ij}(t) + v_{ij}(t+1)$$

PSO with Single Robot

Co-Learning Collaborative Behavior

Three orthogonal axes to consider (extremities or balanced solutions are possible):

- Individual and group fitness
- Private (non-sharing of parameters) and public (parameter sharing) policies
- Homogeneous vs. heterogeneous systems

Private & Individual

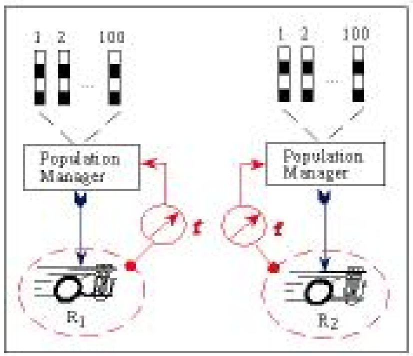

Public & Individual

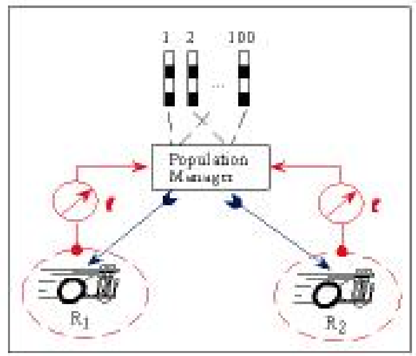

Example with binary encoding of candidate solutions

Private & Group

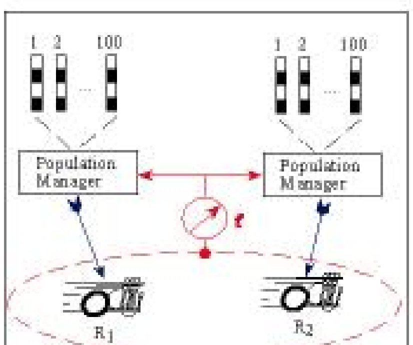

Public & Group (het)

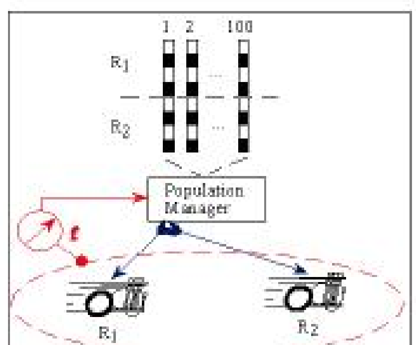

Public & Group (hom)

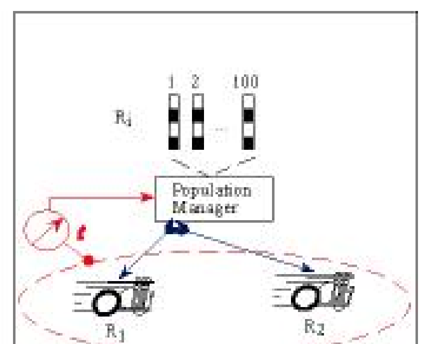

Block V – Selected Topics in Distributed Sensing and Action

- Sensor networks (static, mobile, robotic)
- Self-aggregation and self-assembling

Wireless Sensor Networks

Features:

- Very low sampling frequency < 1Hz
- Very low power consumption: 25mW
- Solar panel
- Radio communication

Sensors:

- Air Temperature and Humidity
- Infrared Surface Temperature
- Anemometer
- Solar Radiation
- Pluviometer
- Soil moisture
- Soil pressure

http://sensorscope.epfl.ch

At DISAL [Evans and Bahr, 2009 - 2015]

OpenSense

Air Pollution Monitoring

NANO

SENSING SYSTEM
From many wireless, mobile, heterogeneous, unreliable raw measurements ...

INFORMATION SYSTEM
... to reliable, understandable and Web-accessible real-time information

TERA

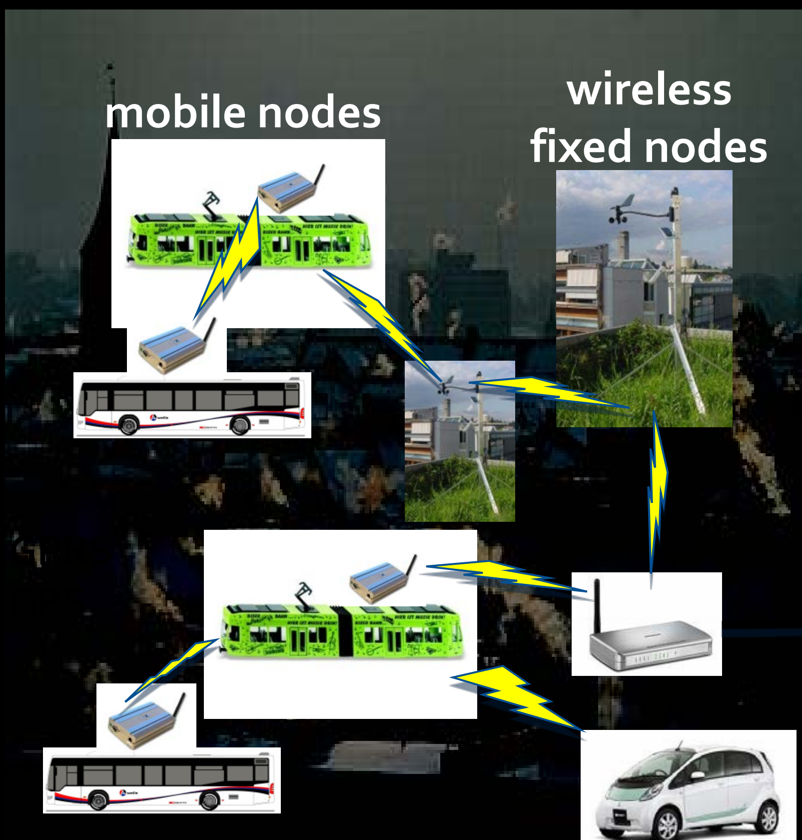

electric vehicle (C-Zero)

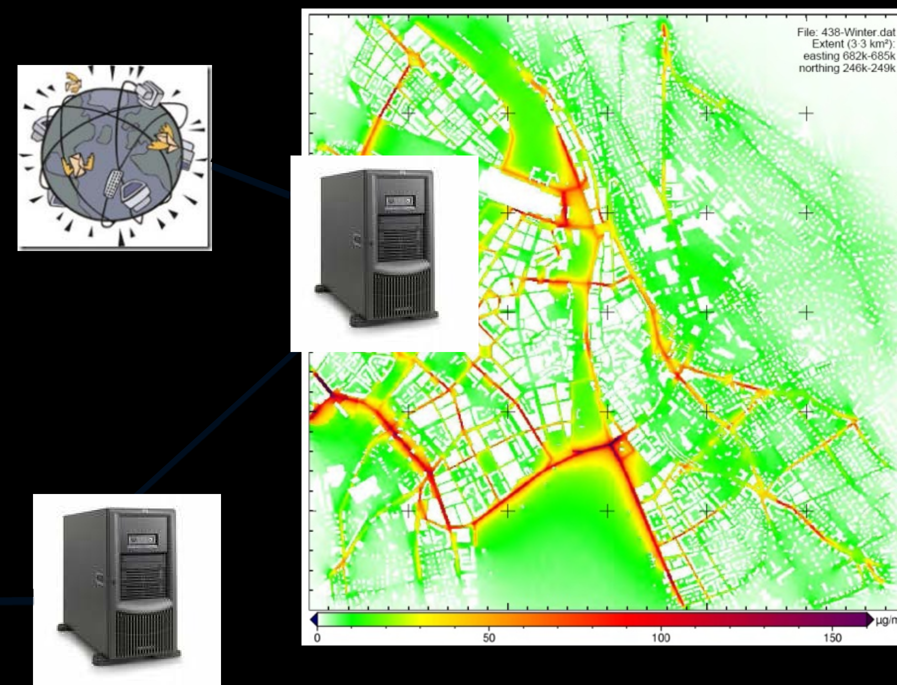

At DISAL: [Arfire, 2010 - 2016; Marjovi 2014 -]

Distributed Odor Source Localization

- Bio-inspired and information theoretical algorithms
- Distributed control and sensing
- Integration of anemometry, olfaction, and inter-robot localization capabilities
- Wind tunnel and simulation experiments
- Possible applications: environmental pollution, search and rescue operations, humanitarian demining
- [Lochmatter, 2005 - 2010]; [Soares 2011-2016]; [Marjovi, 2014 -], Rahbar [2015 -]

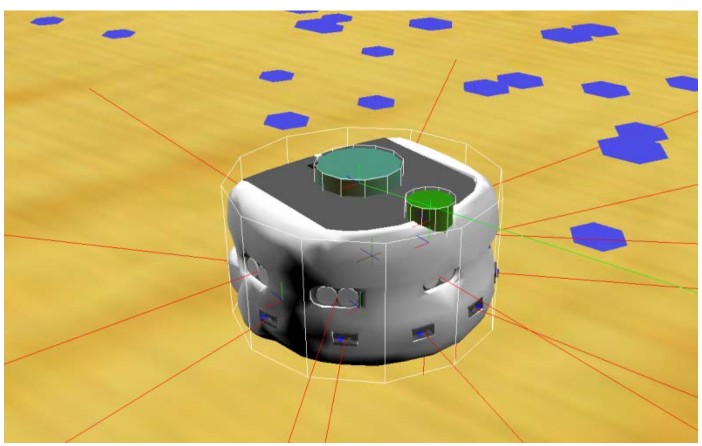

Self-Aggregation

- Given an delimited area containing two similar or dissimilar (i.e. one bright and one dark) shelters infiltrate using mechatronic lures the insect society and participate to the collective decision, voting for a natural or artificial solution
- Applications:
  - Low-stress animal management
  - Alternative pest control
- Robot endowed with:
  - 1D Vision
  - Chemical camouflage
  - IR-based navigation and com modules

[Halloy et al, *Science* 2007]

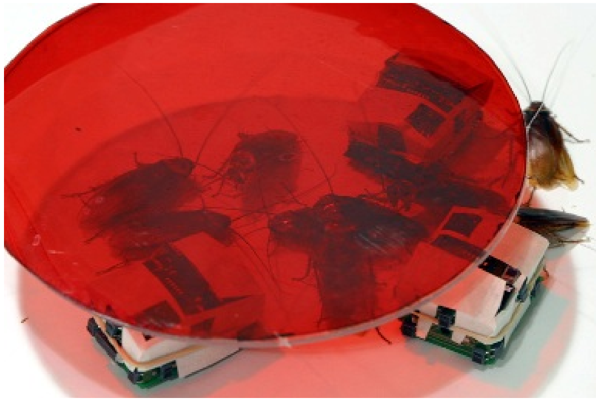

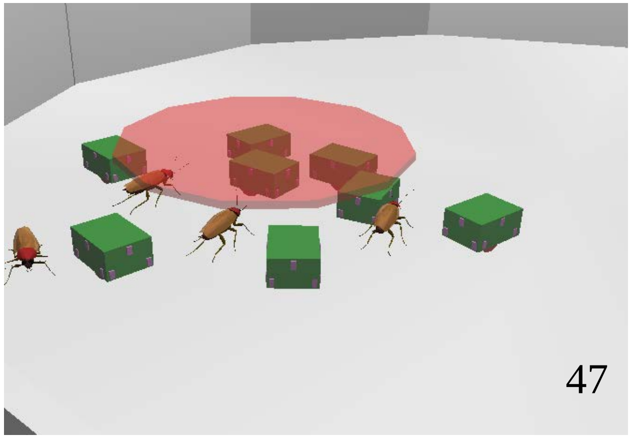

Self-Assembling

Alice mobile robot

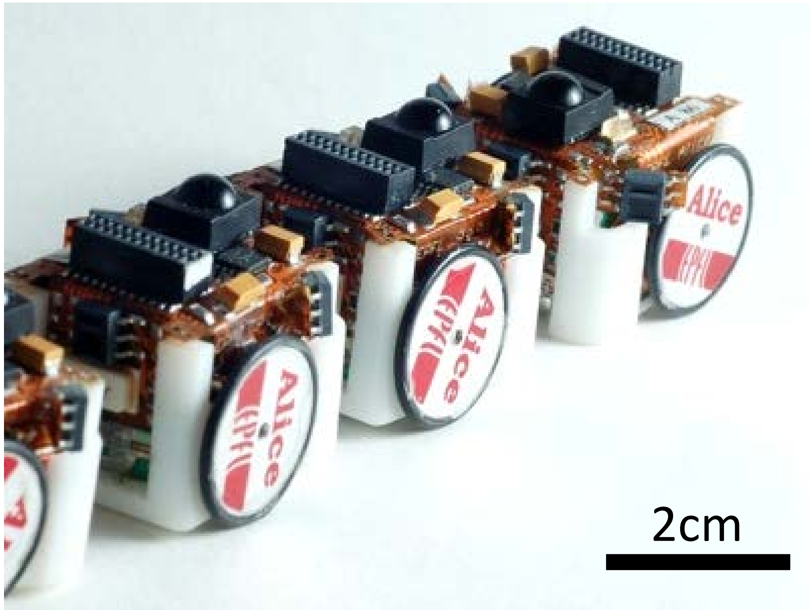

SU-8 microfabricated parts

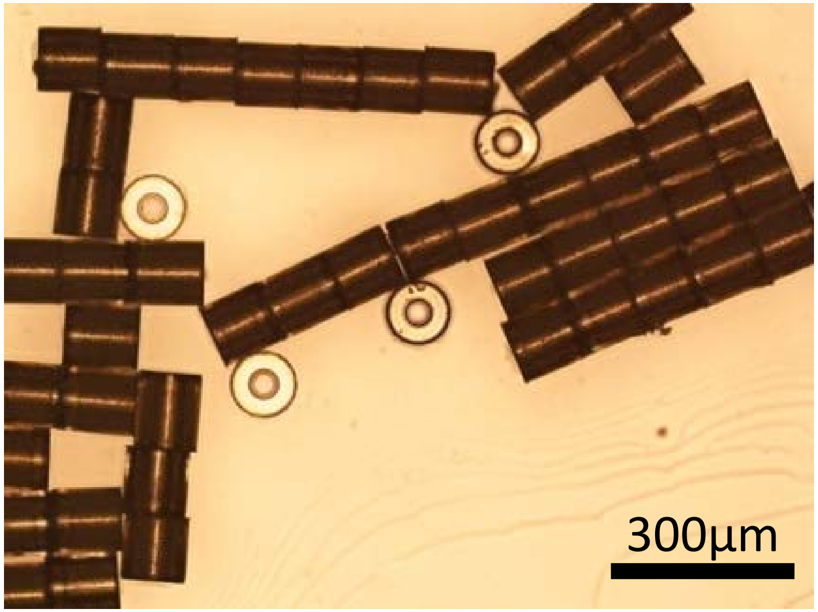

- Size: 2 centimeters
- Typical swarm size: a few dozen units
- Sensing, computation, communication
- Controllable (but noisy) self-locomoted units

- Size: 50 to 500 µm
- Typical swarm size: 10^3 to 10^6 units
- No sensing, no computation, no communication, but **local interactions**
- No self-locomotion, stochastic motion

[Mermoud, 2006 - 2012], [Haghighat 2012 -]

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Conclusion

Take Home Messages

1. Course is rich and intensive; consider your overall semester load before enrolling
2. Balanced theoretical contents and hands-on experience; first two lab sessions give an idea of the workload
3. The course is close to the research in its purpose and remains a showcase of what we do (biased selection of topics and material distributed)
4. Check previous editions on the web (especially last edition because very close in its structure) for exercises, exam questions, discuss with TAs if appropriate, to decide whether to definitively enroll in the course

Distributed Intelligent Systems – W1

Part II: An Introduction to Swarm Intelligence, Foraging Strategies in Ant Societies, and Ant-Inspired Metaheuristics

Outline

- Swarm Intelligence
  - A possible paradigm and motivation
  - Key principles
- Foraging Strategies
  - Recruitment-based mechanisms
  - Inaccuracies of chemical communication
- Bridges experiments in the lab
- Open space and multi-source experiments
- Ant networks
- The Traveling Salesman Problem (TSP)
- An Ant System (AS) for the TSP

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

An Introduction to Swarm Intelligence – Motivation, Definitions, and Key Principles

Some natural collective phenomena implying a close interconnection among individuals

...

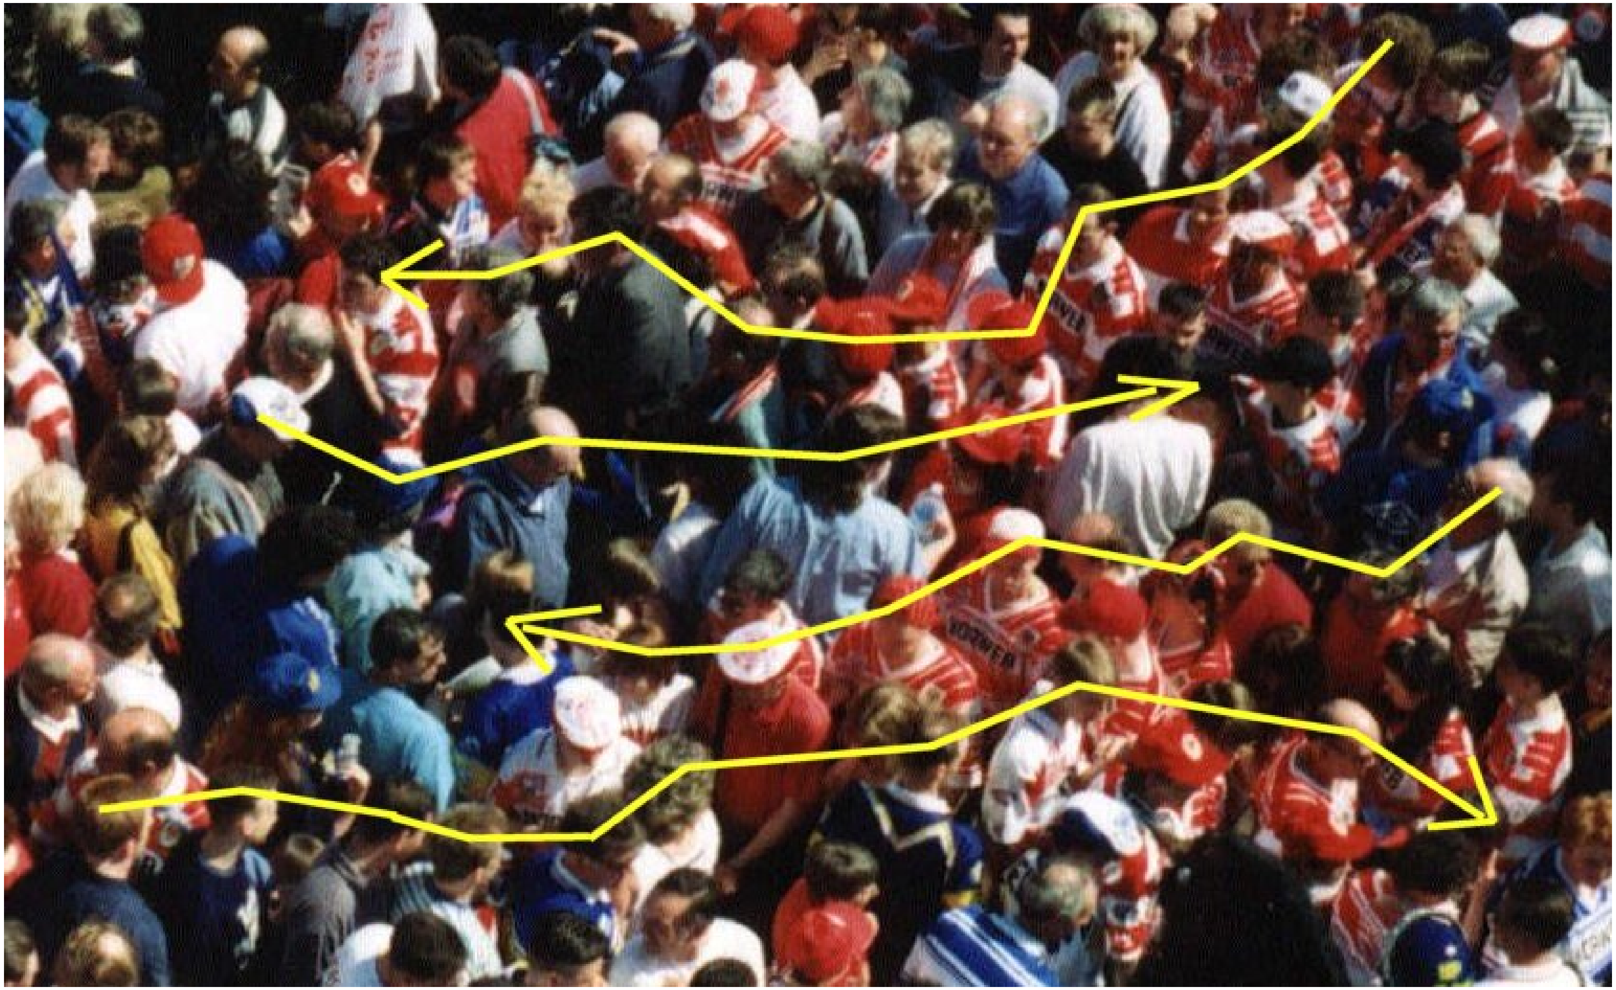

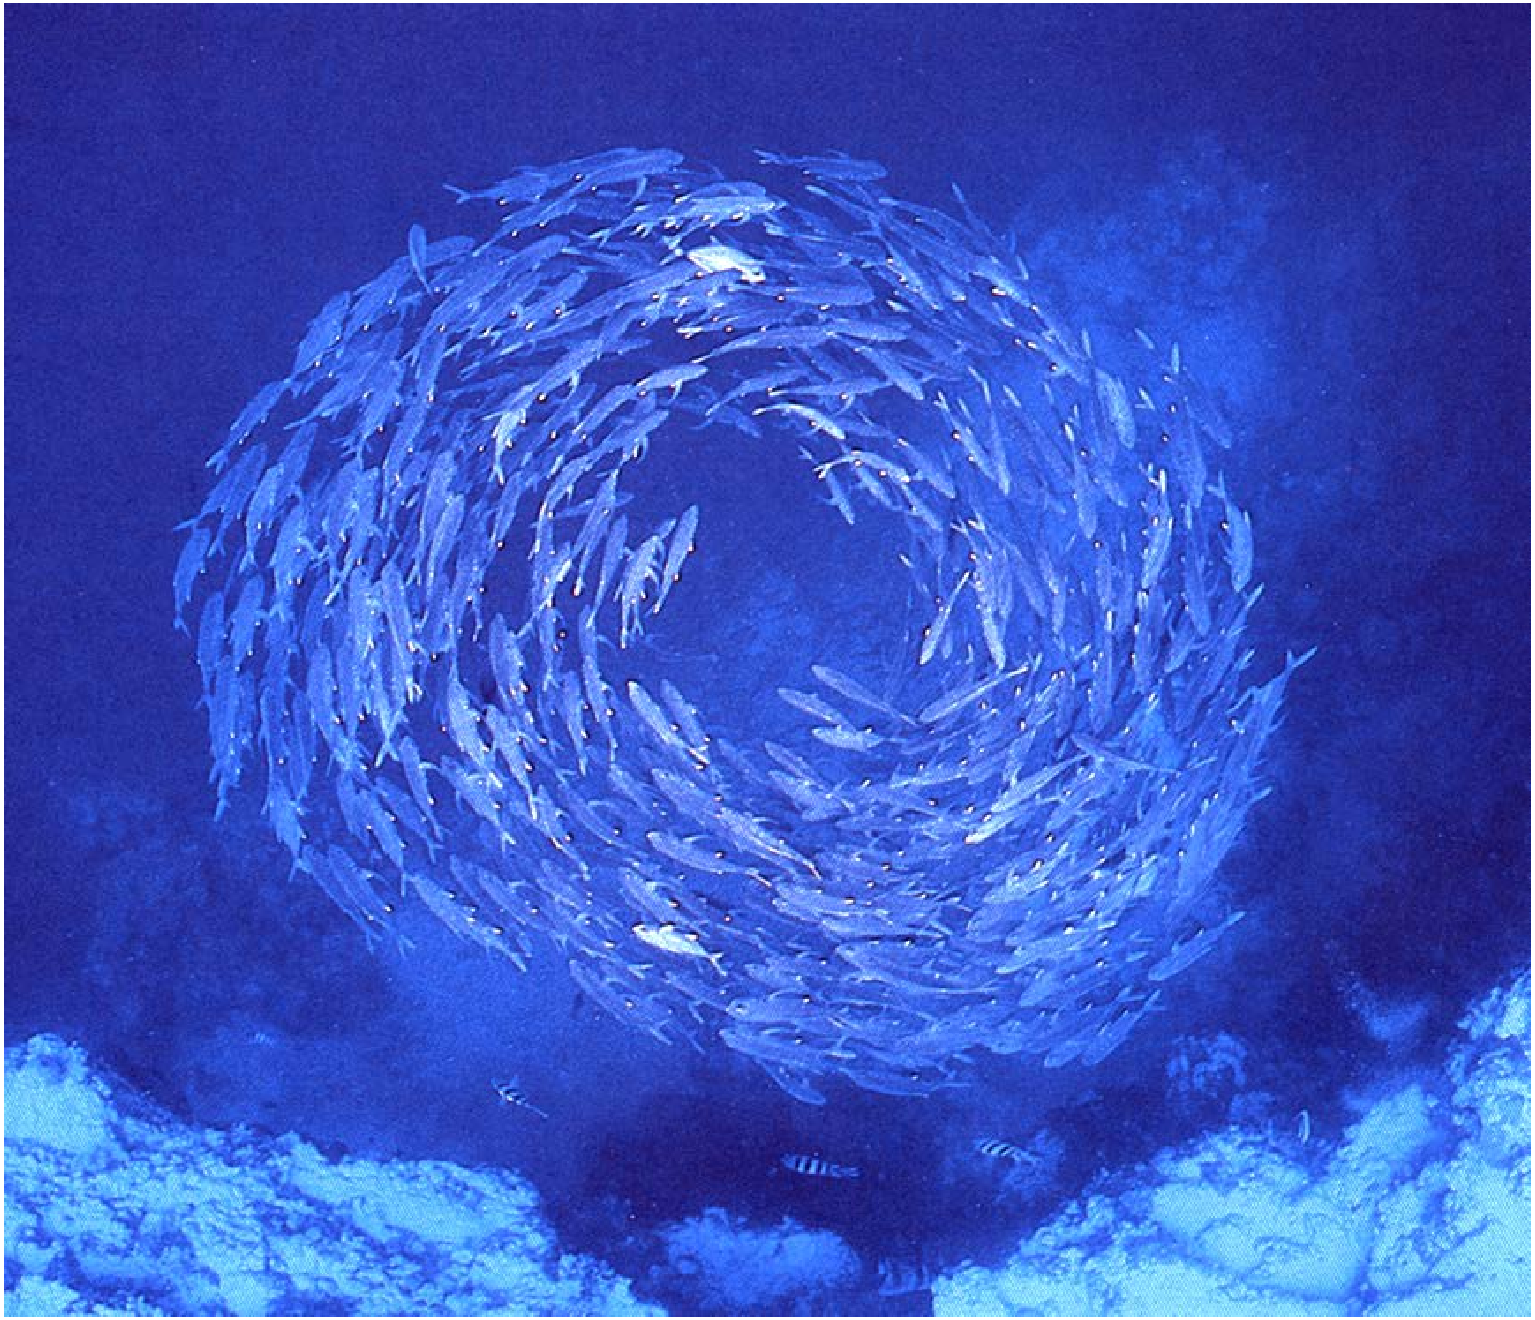

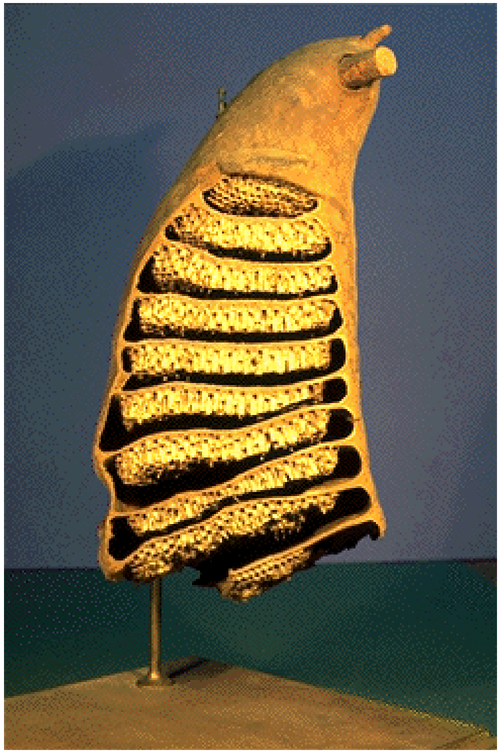

© Guy Theraulaz, UPS, 1999

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Collective Phenomena

- **Limited local information**

  Each individual in the group **has access only to limited local information and has no global knowledge of the structure** which it is engaged in constructing together with the other members of the group

- **A set of simple individual rules**

  Each individual obeys **a collection of a few simple behavioral rules**. This rule set permits the group collectively to coordinate its activities and to build a global structure or configuration.

- **The global structures which emerge accomplish some function**

  These structures often allow the group to solve problems. They are **flexible** (adapting easily to a novel environment), and they are **robust**, (if one or several individuals fail in their behavior or make a simple mistake, the structures spontaneously re-form).

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

From Natural to Artificial Systems and more ...

Collective/Swarm Intelligence?

Some questions arise ...

- How do animal societies manage to perform difficult tasks, in dynamic and varied environments, without any external guidance or control, and without central coordination?

- How can a large number of entities with only partial information about their environment solve problems?

- How can collective cognitive capacities emerge from individuals with limited cognitive capacities?

From Natural to Artificial Systems

- Modeling to understand microscopic to macroscopic transformation
- Modeling as interface to artificial systems

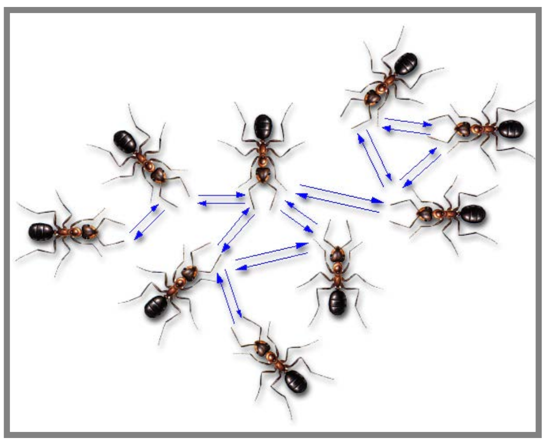

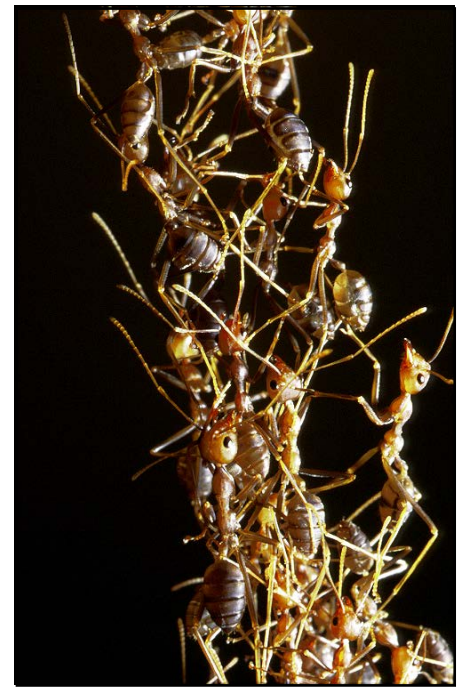

Individual behaviors and local interactions

Modeling

Ideas for artificial systems

Global structures and collective decisions

Computational Swarm-Intelligent Systems

- In a virtual world, most of the mechanisms shown by natural SI can be easily reproduced
- Some of the mechanisms are intentionally modified and further ones are added in order to improve the performance of a given algorithm

Physical Swarm-Intelligent Systems

- Beyond bio-inspiration: combine natural principles with engineering knowledge and technologies
- Unit coordination
  - fully distributed control (+ env. template)
  - individual autonomy
  - self-organization (extend definition)
- Communication
  - explicit/implicit local communication
  - indirect communication through signs in the environment (stigmergy)
- Scalability
- Robustness vs. efficiency trade-off
  - redundancy
  - balance exploitation/exploration
  - individual simplicity
- System cost effectiveness
  - individual simplicity
  - mass production

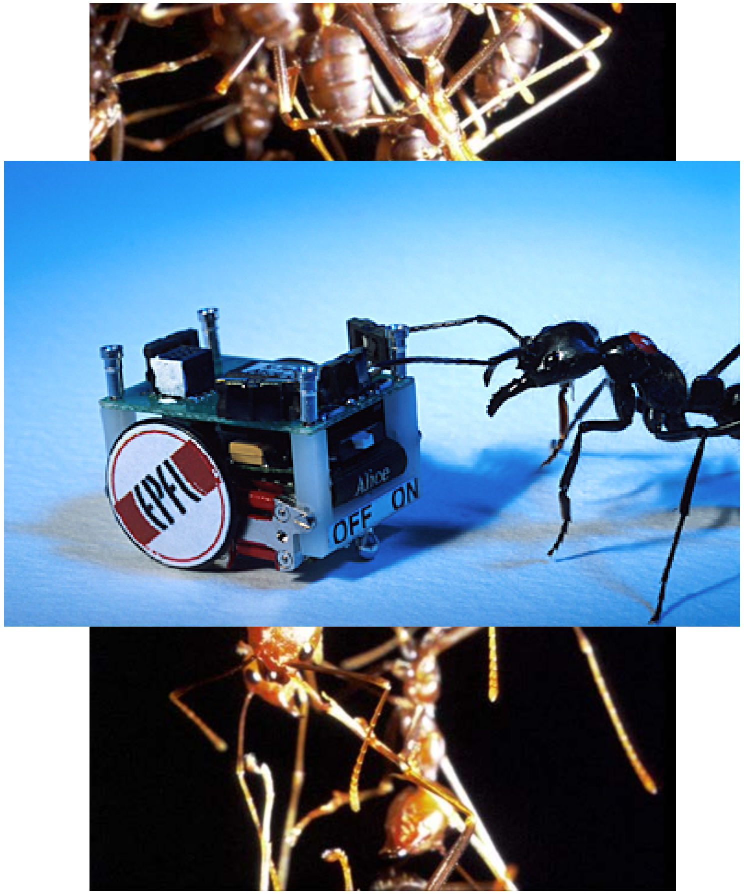

Some Definitions of Swarm Intelligence

- Beni and Wang (1989):
  - Used the term in the context of cellular automata (based on cellular robots concept of Fukuda)
  - Decentralized control, lack of synchronicity, simple and (quasi) identical members, self-organization
- Bonabeau, Dorigo and Theraulaz (1999)
  - Any attempt to design algorithms or distributed solving devices inspired by the collective behavior of social insect colonies and other animal societies
- Beni (2004)
  - Intelligent swarm = a group of non-intelligent robots ("machines") capable of universal computation
  - Usual difficulties in defining the "intelligence" concept (non predictable order from disorder, creativity)

Key Mechanisms behind Natural Swarm Intelligence

Two Key Mechanisms in Natural Swarm-Intelligent Systems

1. Self-Organization

2. Stigmergy

Self-Organization

- Set of dynamical mechanisms whereby **structure appears at the global level** as the result of **interactions among lower-level components**

- The rules specifying the interactions among the system's constituent units are executed on the basis of **purely local information**, without reference to the global pattern, which is an **emergent property of the system** rather than a property imposed upon the system by an external ordering influence

Characteristics of Natural Self-Organized Systems

- Creation of spatio-temporal structures
  - E.g., foraging trails, nest architectures, clusters of objects, ...
- Multistability
  (i.e., possible co-existence of several stable states)
  - E.g., ants exploit only one of two identical food sources, build a cluster in one of the many possible locations, ...
- Existence of bifurcations when some parameters change
  - E.g., termites move from a non-coordinated to a coordinated phase only if their density is higher than a threshold value

Basic Ingredients of Natural Self-Organized Systems

- **Multiple interactions**
- **Randomness**
- **Positive feedback**
  - E.g., recruitment, reinforcement
- **Negative feedback**
  - E.g., limited number of available foragers, pheromone evaporation

Stigmergy

Grassé P. P., 1959

- “La coordination des taches, la regulation des constructions ne dependent pas directement des oeuvriers, mais des constructions elles-memes. ***L’ouvrier ne dirige pas son travail, il est guidé par lui***. C’est à cette stimulation d’un type particulier que nous donnons le nom du **STIGMERGIE** (*stigma*, piqure; *ergon*, travail, oeuvre = oeuvre stimulante).”

- [“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***. It is to this special form of stimulation that we give the name **STIGMERGY** (*stigma*, sting; *ergon*, work, product of labor = stimulating product of labor).”]

Stigmergy

Definition

It defines a class of mechanisms exploited by social insects to coordinate and control their activity via indirect interactions.

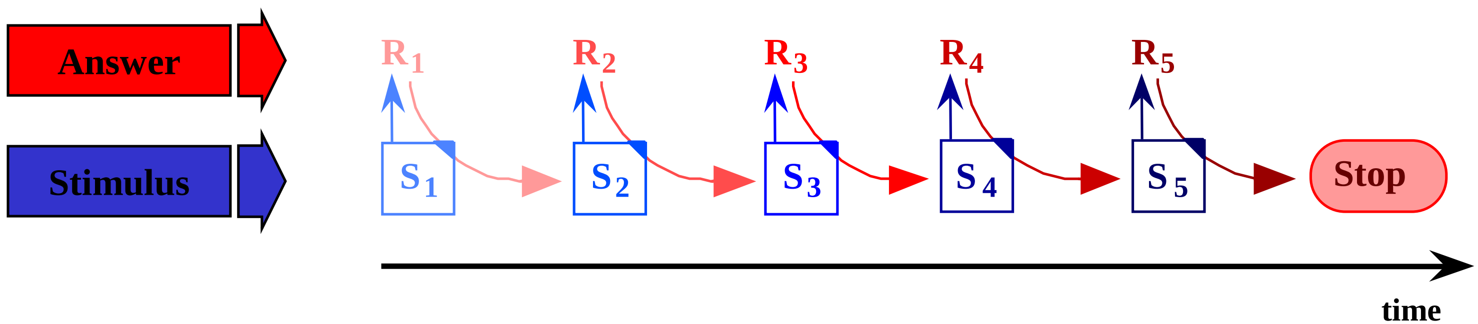

Stigmergic mechanisms can be classified in two different categories:

- quantitative (or continuous) stigmergy
- qualitative (or discrete) stigmergy

[Theraulaz & Bonabeau., *Alife J.* 1999]

Foraging Strategies in Ants

Different Ants, Different Strategies

Not All Foraging Strategies are Collective and based on Stigmergy …

- Example: Cataglyphis desert ant
- Excellent study by Prof. R. Wehner (University of Zuerich, Emeritus)
- Individual foraging strategy
- Underlying mechanisms
  - Internal compass (polarization of sun light)
  - Dead-reckoning (path integration on neural chains for leg control)
  - Local search (around 1-2 m from the nest)
- Extremely accurate navigation: averaged error of a few tens of cm over 500 m path!

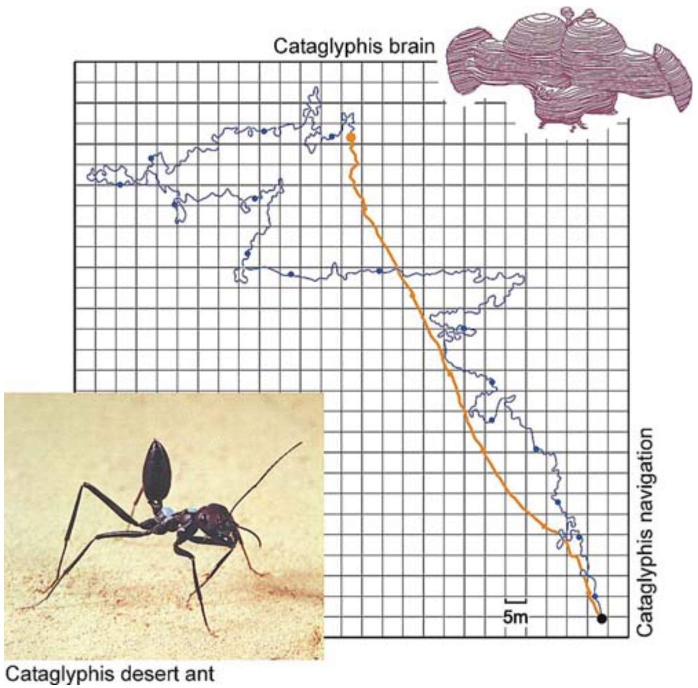

More individual Foraging Strategies

Individual navigation + learning capabilities for memorizing the foraging zone

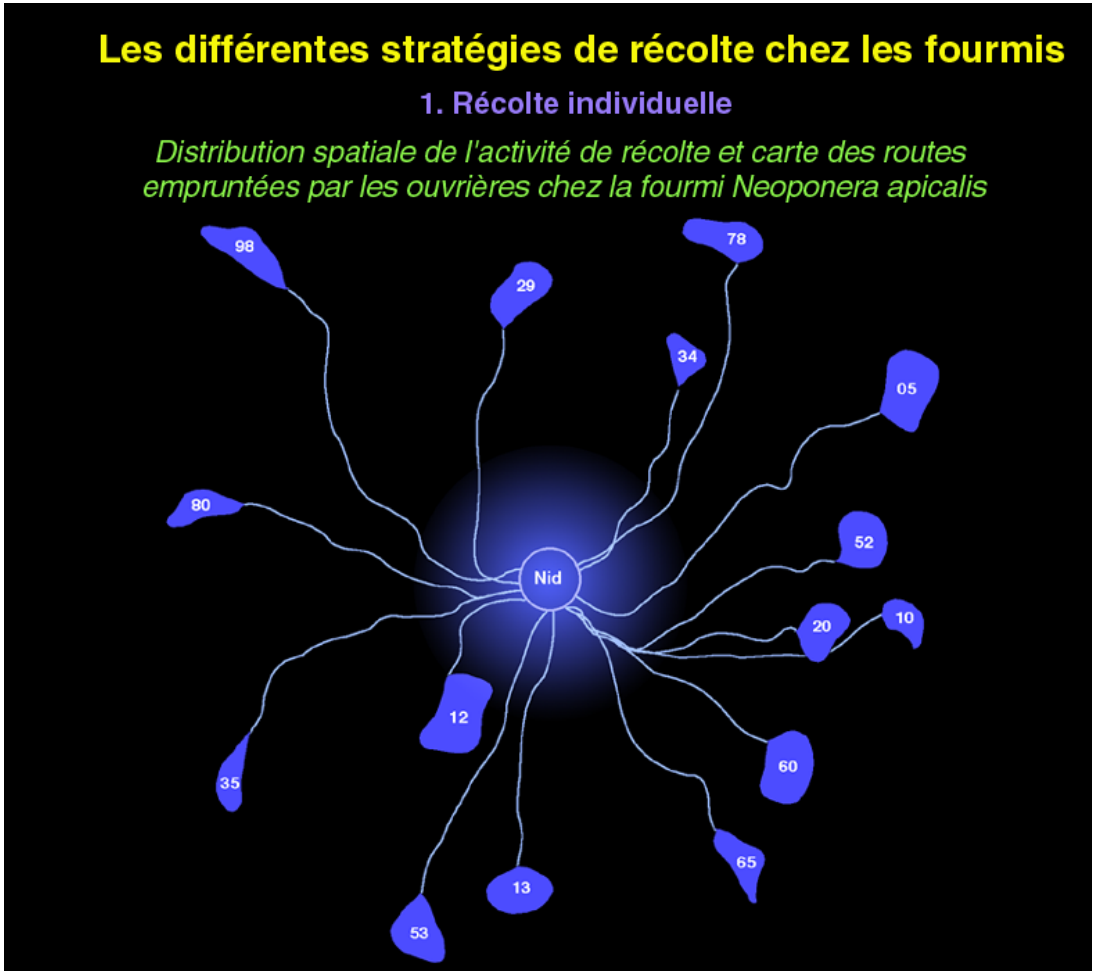

Tandem Recruitment Strategies

- Mediated by thropallaxis, antennal contact
- Based on food chemical signatures on the ant body

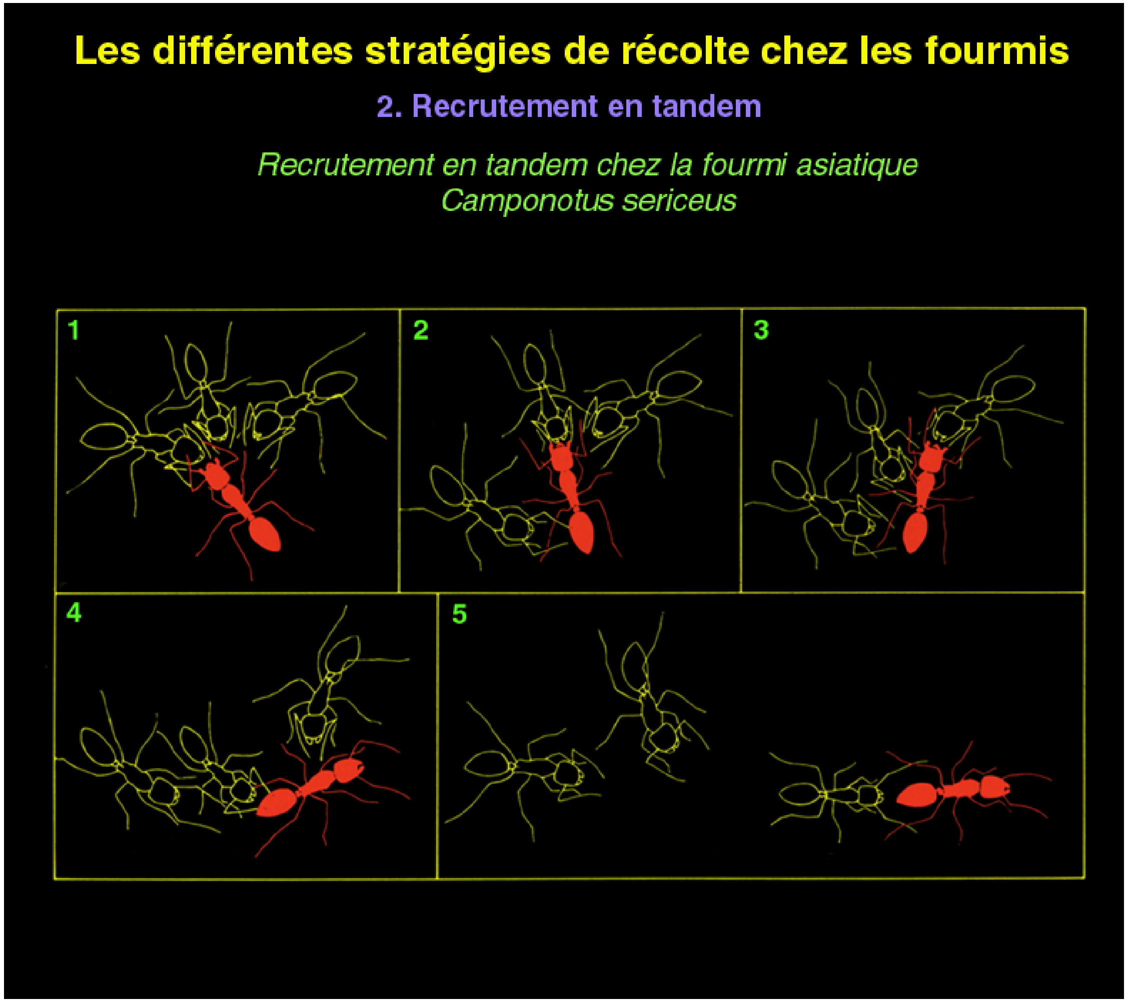

Group Recruitment Strategies

Mass Recruitment Strategies

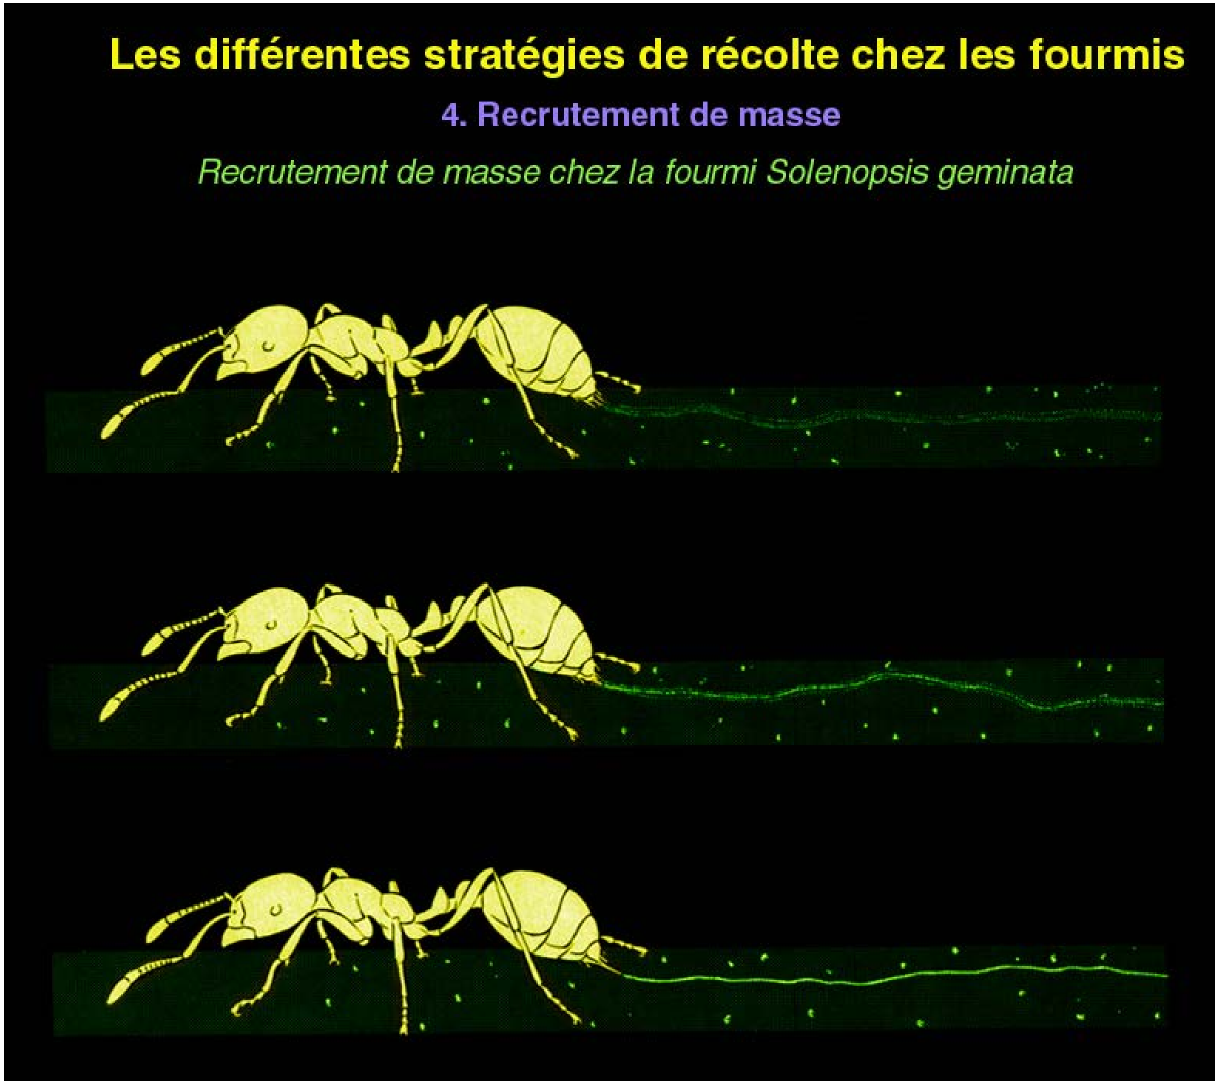

Mass Recruitment

Behavior of Individual Ants

Sequence of actions performed by an ant communicating the discovery of a food source

Food source	Foraging area	Nest

Stimulating nest mates

Laying a chemical trail

Picking up food

Deposition of food

Following the trail

Formation of Recruitment Trails in Ants

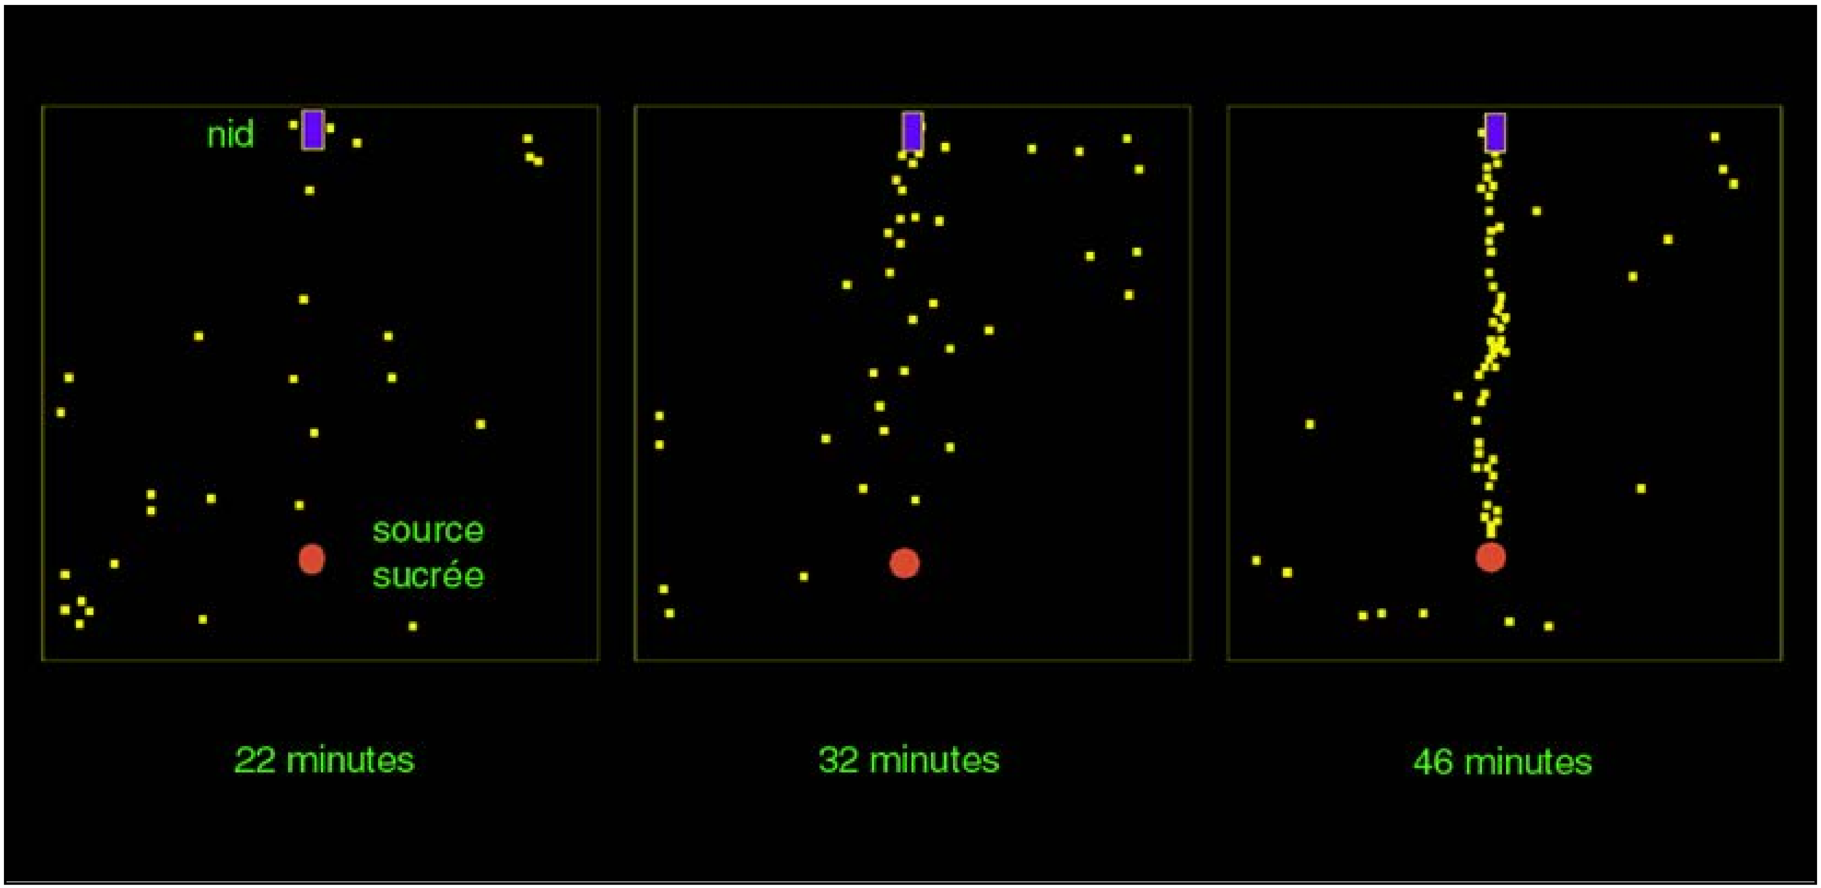

Number of Ants at the Food Source vs. Time

Growing phase (positive feedback)

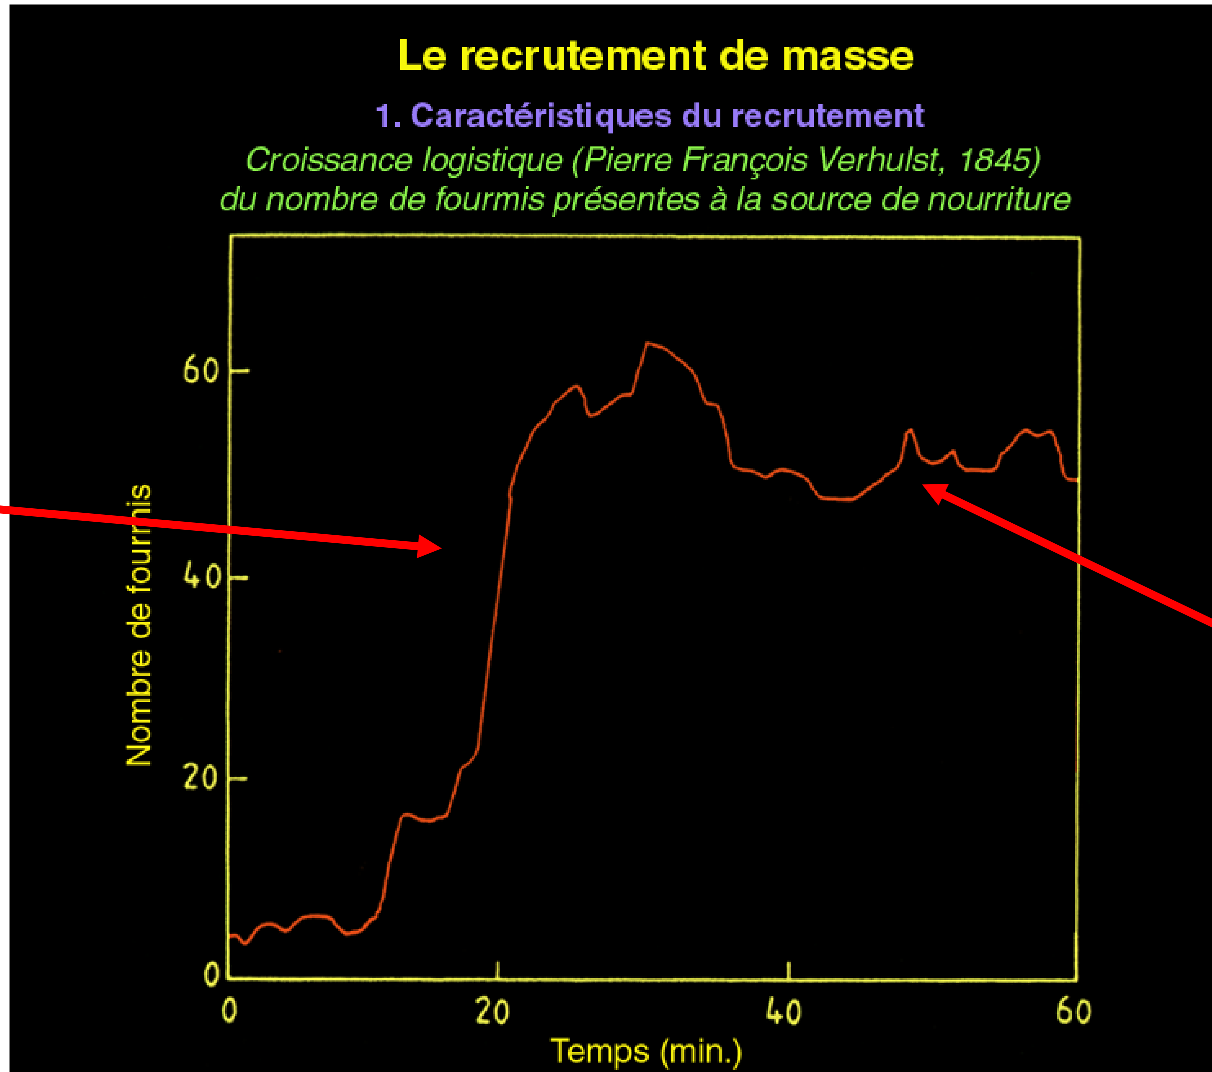

Saturation phase (negative feedback)

Stochastic Individual Behavior Combined with the Amplification of Information can lead to Collective Decisions

The Role of Randomness in the Organization of Foraging

How does individual behavior with a strong stochastic component lead to **statistically predictable behavior** at the level of the colony and **collective decisions**?

Experimental Strategy

- Most of the studies to assess quantitatively the role of randomness have been carried out in the lab because:
  - Controlled environmental conditions
  - Repeated runs for statistics

- Studies in the field can lead often only to qualitative conclusions because they might be influenced by:
  - Multiple food sources
  - Predators and competitors
  - Environmental changes (temperature, climate, etc.)

Exploration: The Inaccuracy of Chemical Communication

Termite Following a Pheromone Trace

Prof. J.-L. Deneubourg (ULB, Bruxelles)

Ants can Reacquire a Trail by Local Search

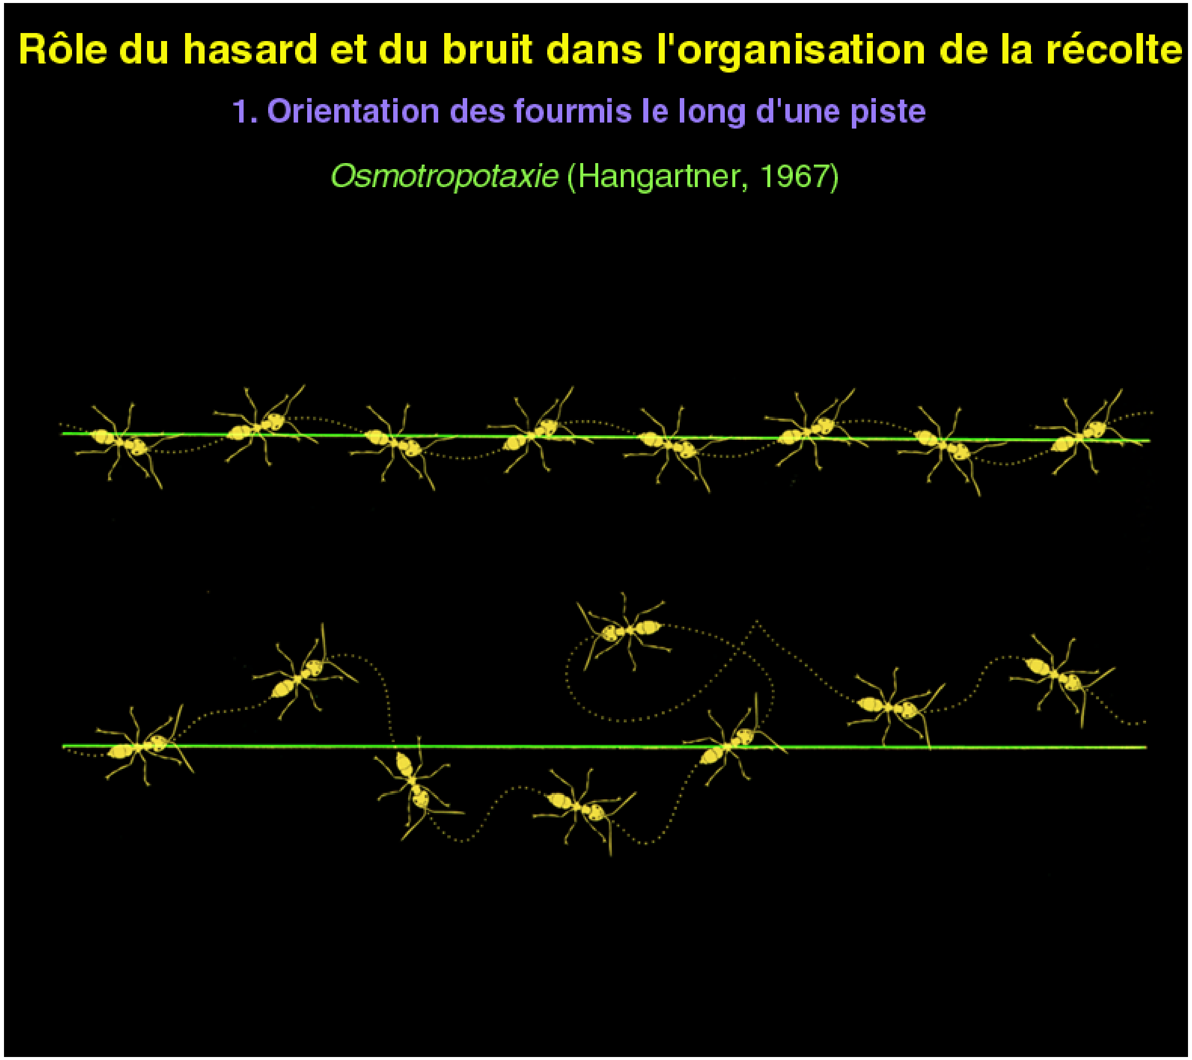

Probability of Trail Losing depends on the Ant Species

Example: Accuracy of recruitment of the first recruit (Verhaeghe et al., 1980)

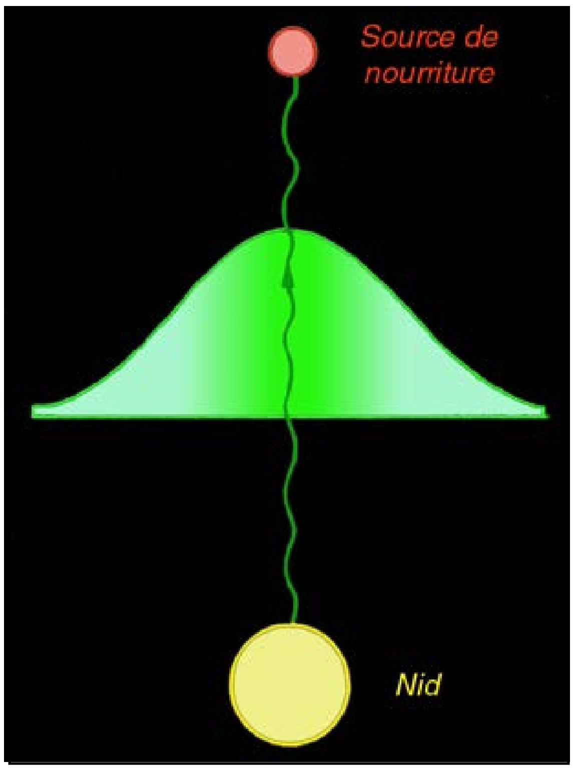

	Tetramorium impurum	***Tapinoma erraticum***
Successful recruitments (%)	18	74
Length of trail followed (%)	17	68

Tapinoma follow trails much more reliably than *Tetramorium* → depends on the environment the species have evolved (food scattering, etc.)

Probability of Trail Losing is Constant over Time

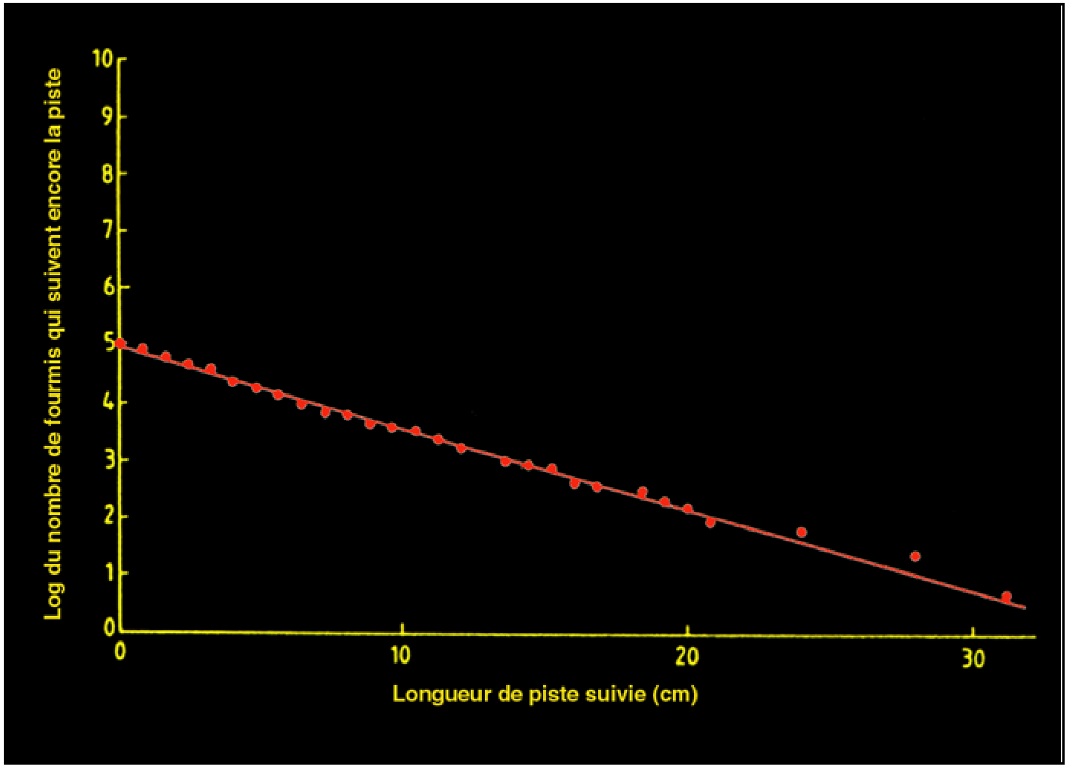

Log # of ants on the trail as a function of the traveled path for a constant pheromone concentration

- The longer the traveled path and the smaller is the number of ants on the trail
- Appears to be independent of phenomena such as learning or sensory adaptive response (at least under such short time scale)

Probability of Trail Losing depends on Chemical Concentration

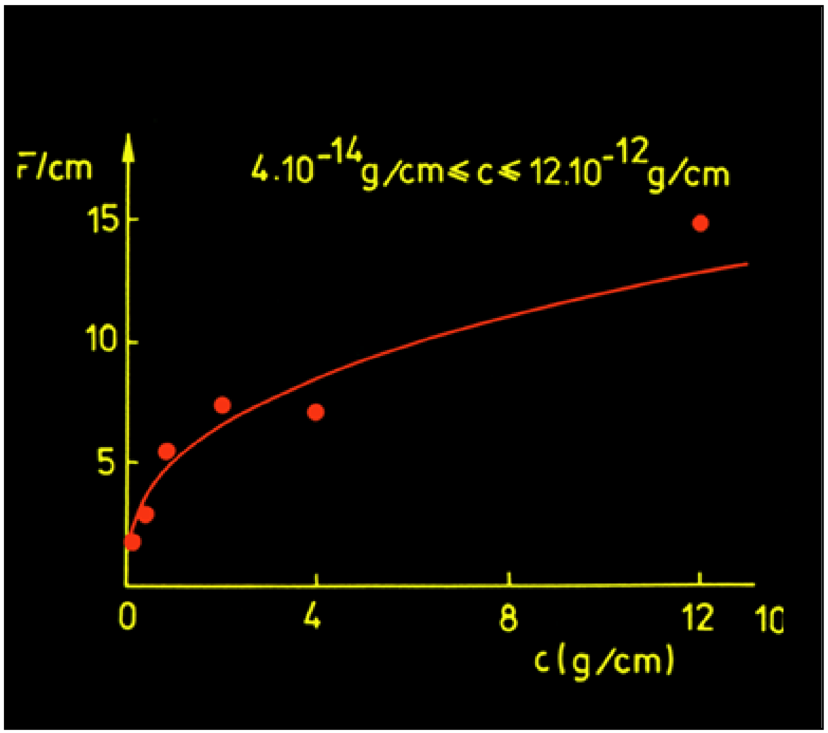

Mean path length as a function of the pheromone concentration

The higher is the pheromone concentration and the more reliably can be followed a trail

Biological Significance of the Exploitation-Exploration Balance

Does the accuracy of the chemical communication channel used by ants increase or decrease their efficiency?

- Noise can have a certain flexible value for the organization of the society.
- The fact that a significant proportion of recruits get lost en route can be of benefit when food is scattered throughout the environment or when several sources are present simultaneously.
- If too many ants get lost for a given food scattering the efficiency of recruitment also decreases.

Sacrifice a little bit efficiency in order to be robust at facing environmental unpredictability

Bridge Experiments: Selecting the Shortest Path

The Suspended, Symmetric Bridge Experiment

Food source

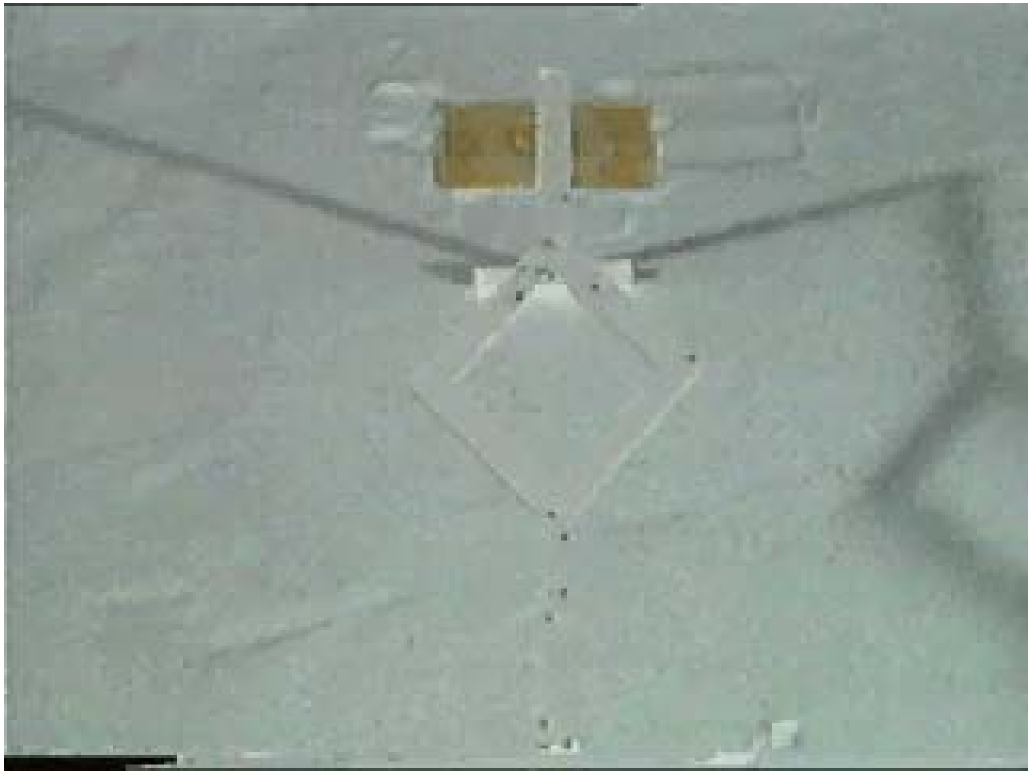

Nest

Two branches (A and B) of the same length connect nest and food source

Experimental Results

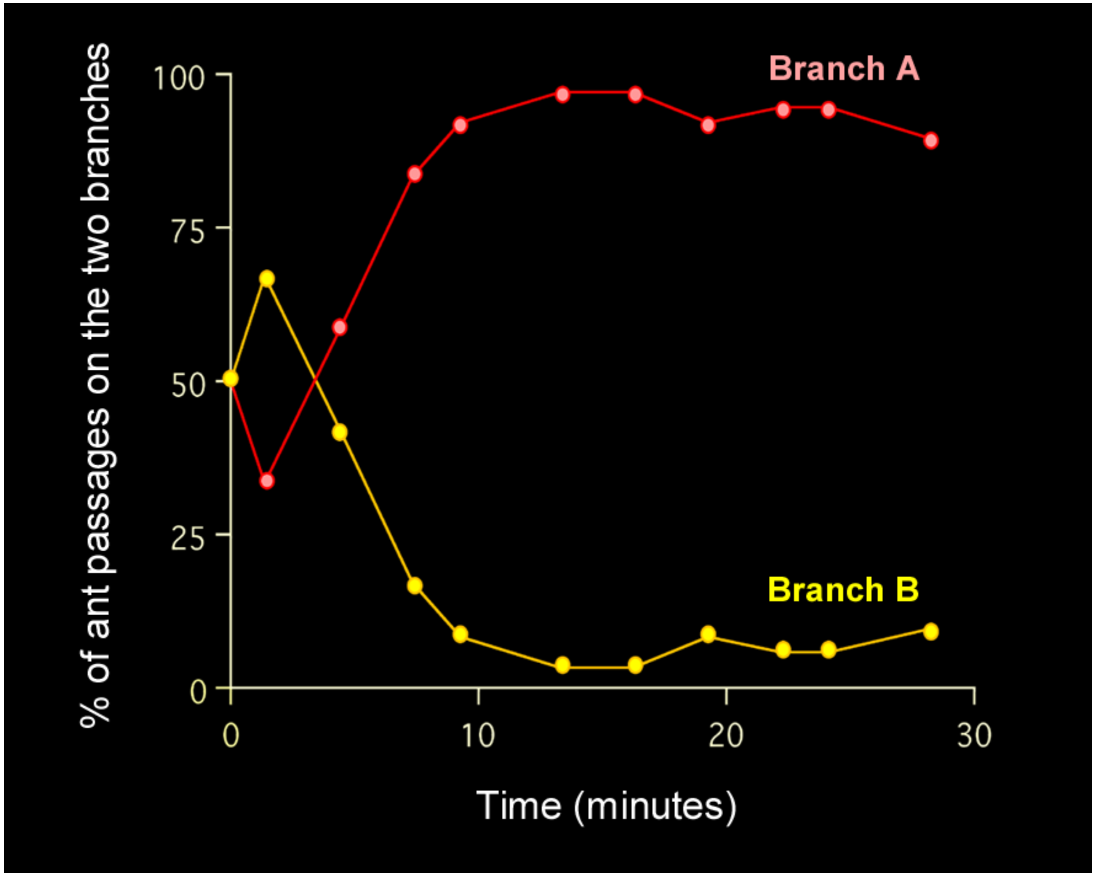

Microscopic Model (Deneubourg 1990)

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

Probabilistic choice of an agent at the bridge's bifurcation points

$\mathbf{P_A}$ and $\mathbf{P_B}$: probability for the ant i+1 to pick up the branch A or B respectively

$\mathbf{A_i}$: number of ants having chosen branch A

$\mathbf{B_i}$: number of ants having chosen branch B

n (model parameter): degree of nonlinearity
k (model parameter): degree of attraction of a unmarked branch

$$A_{i+1} = \left[\begin{array}{lll} A_i + 1 & if & \delta \le P_A \\ A_i & if & \delta > P_A \end{array} \right. \qquad B_{i+1} = \left[\begin{array}{lll} B_i + 1 & if & \delta > P_A \\ B_i & if & \delta \le P_A \end{array} \right.$$

$$A_i + B_i = i$$

δ = uniform random variable on [0,1]

Parameters of the Choice Function

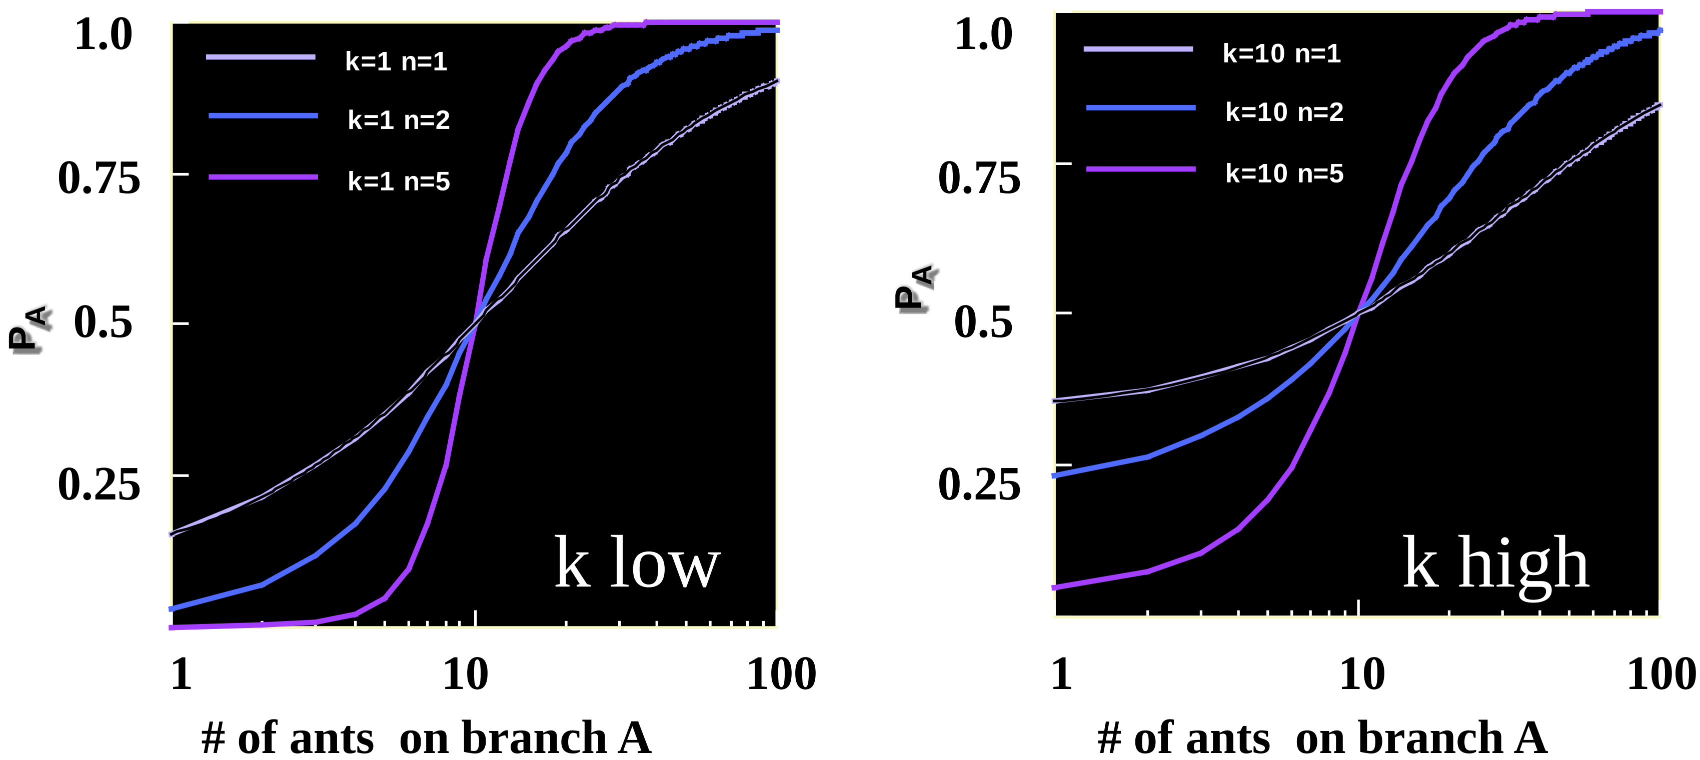

- The higher is n and the faster is the selection of one of the branches (sharper curve); n high corresponds to high exploitation
- The greater k, the higher the attractivity of a unmarked branch and therefore the higher is the probability of agents of making random choices (i.e. not based on pheromones concentration deposited by other ants); k high corresponds to high exploration

Model vs. Experiments

% of ant passages on the dominant branch

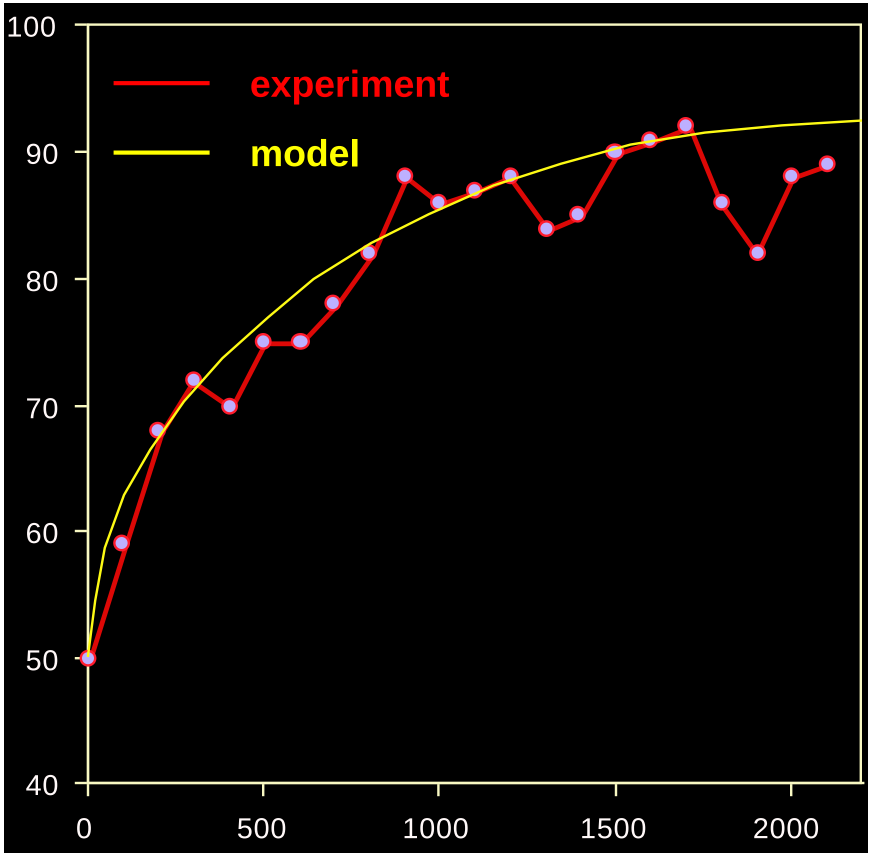

Total number of ants having traversed the bridge

Parameters that fit experimental data:

$n = 2$

$k = 20$

The Suspended, Asymmetric Bridge Experiment

Food source

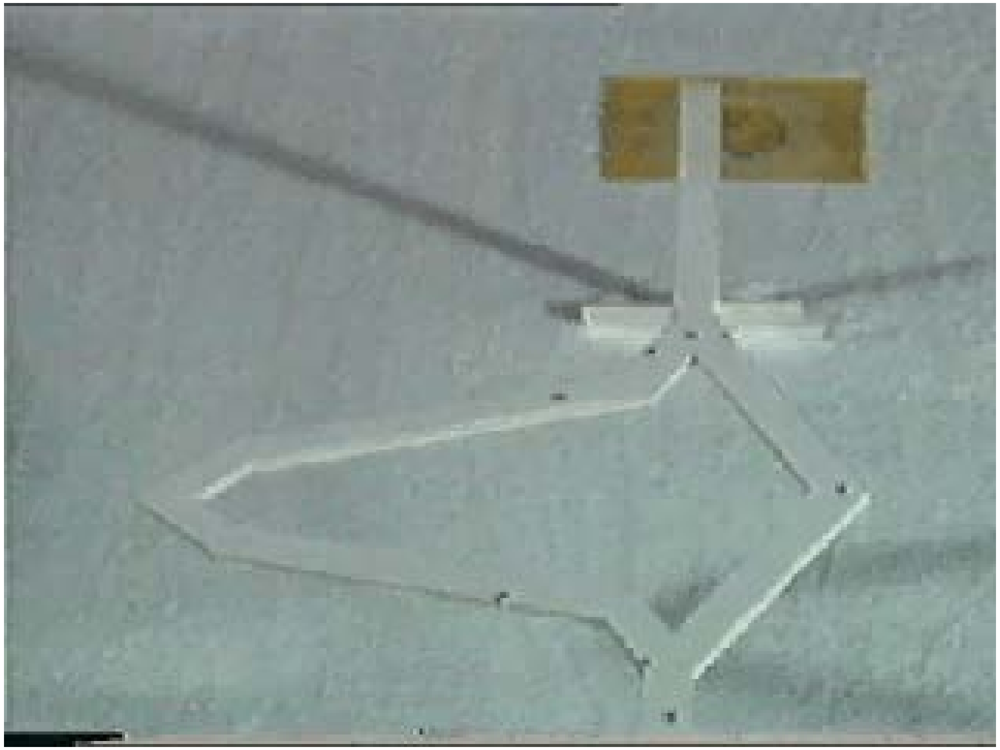

Nest

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

© J.-L. Deneubourg

All Bridge Experiments

4 different experimental scenarios

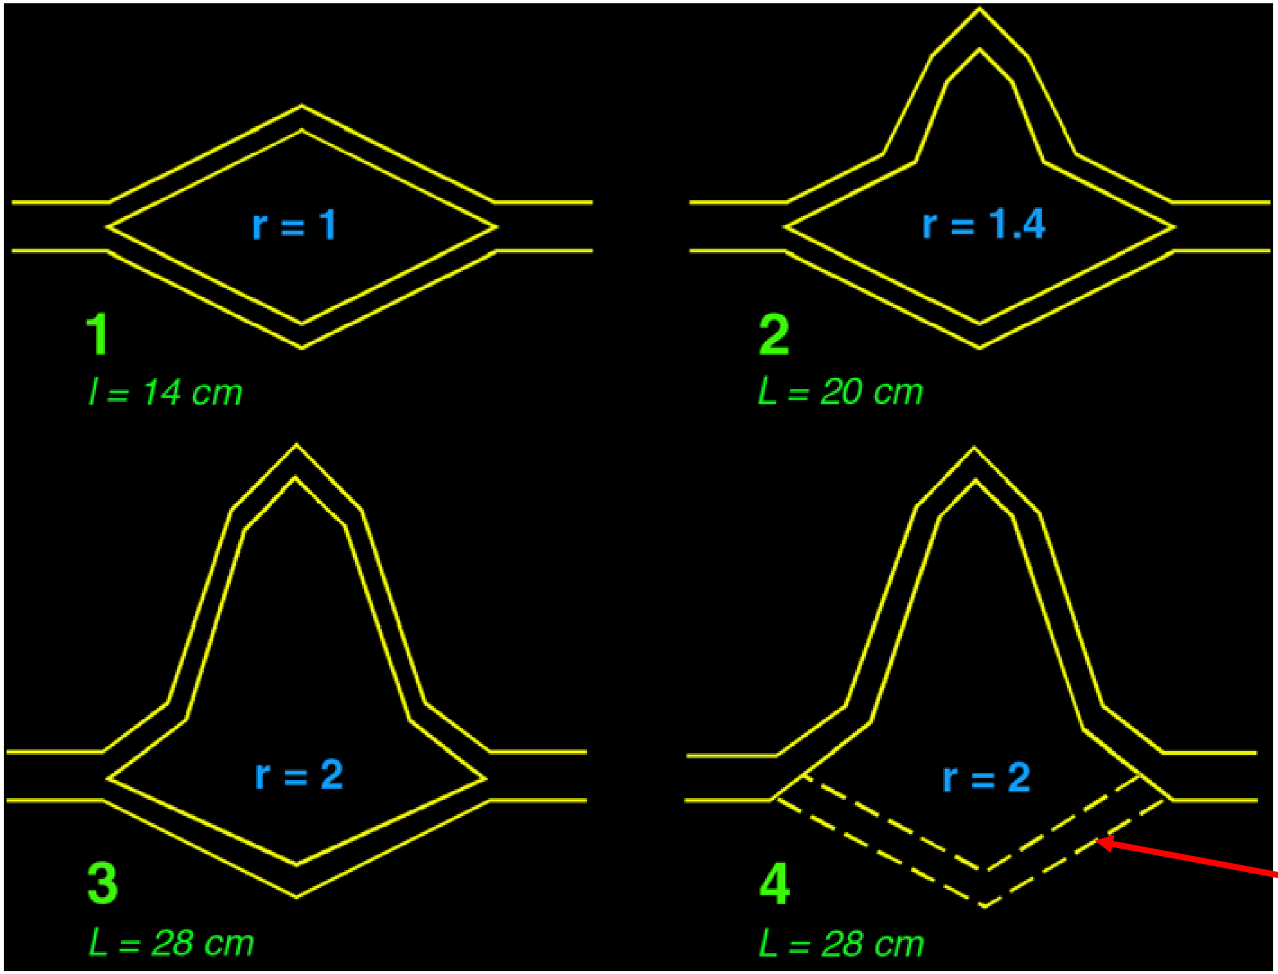

Shortest branch added later

Selection of the Shortest Branch

Repeated experiments with different ant colonies of the same ant species (*Linepithema Humile*) – finite experimental time window

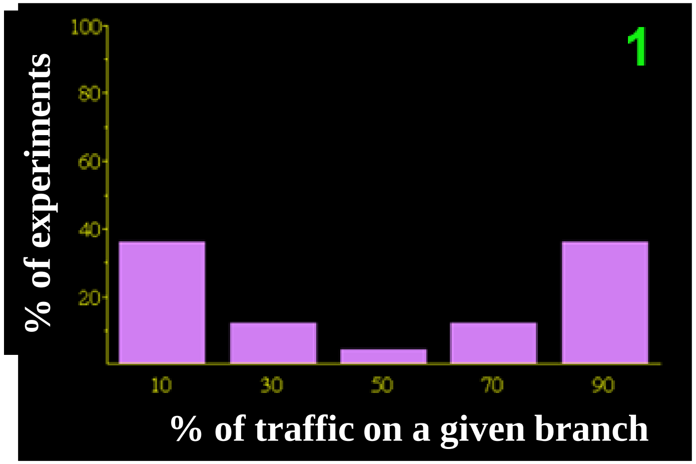

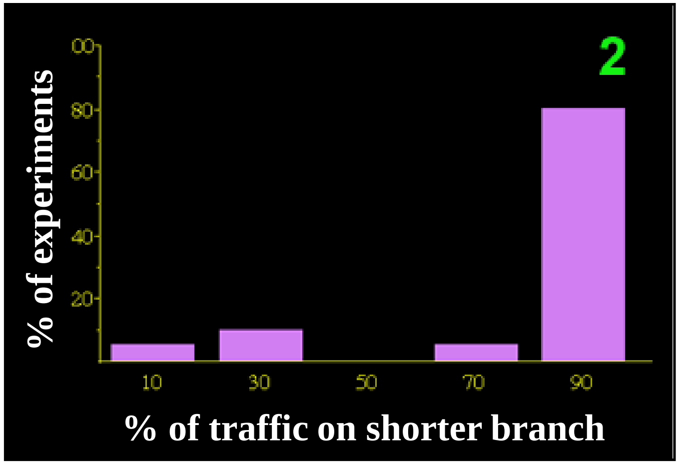

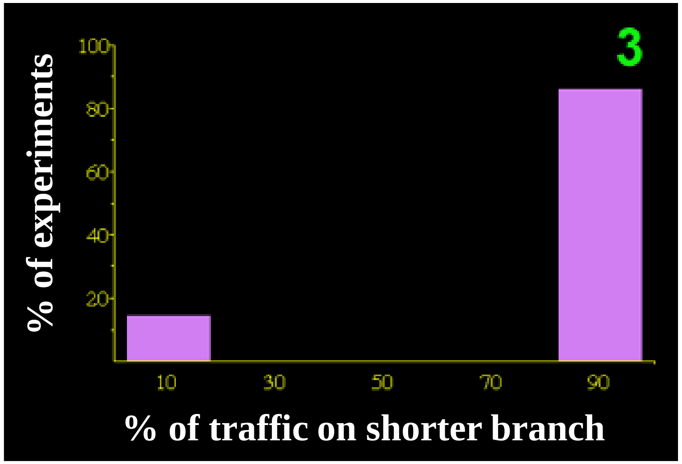

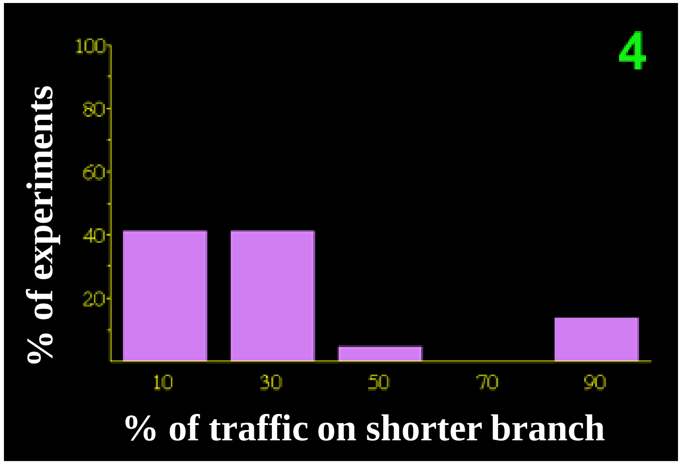

Asymmetric Bridge – Ant Species Differences

What happens if the shorter branch is presented after 30 minutes?

- Argentine Ants (*Linepithema Humile*) get stuck on the longer branch (mainly pheromone-based navigation), see previous slide.
- *Lasius Niger* ants find the shorter branch because they integrate other navigation modalities (compass, dead-reckoning) with pheromone navigation -> U-Turns (different from random walk)!
- *Pharaoh ants* recognize the right way to go from geometry of trails (trails geometry provide polarity information!), again dead-reckoning/compass capabilities!

Asymmetric Bridge – Microscopic Modeling

- The previous model does not work any more: distance/traveling time has to be considered in order to incorporate the geometry of the bridge.
- Multi-agent simulation incorporating pheromone deposition, avoidance rules, … point simulator (take into account trajectories but no body) by Prof. M. Dorigo (ULB Bruxelles).

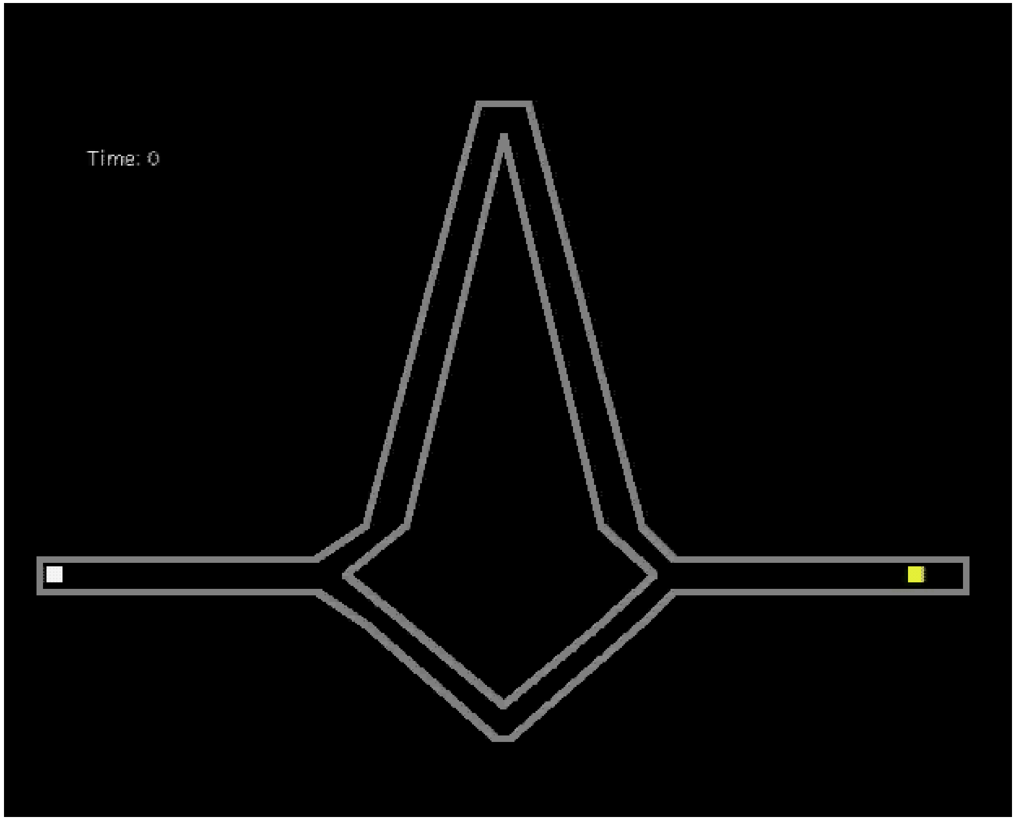

Foraging in Free Space

Selecting the Richest Source

Three different experimental scenarios:

Experiment N°1

Experiment N°2

Experiment N°3

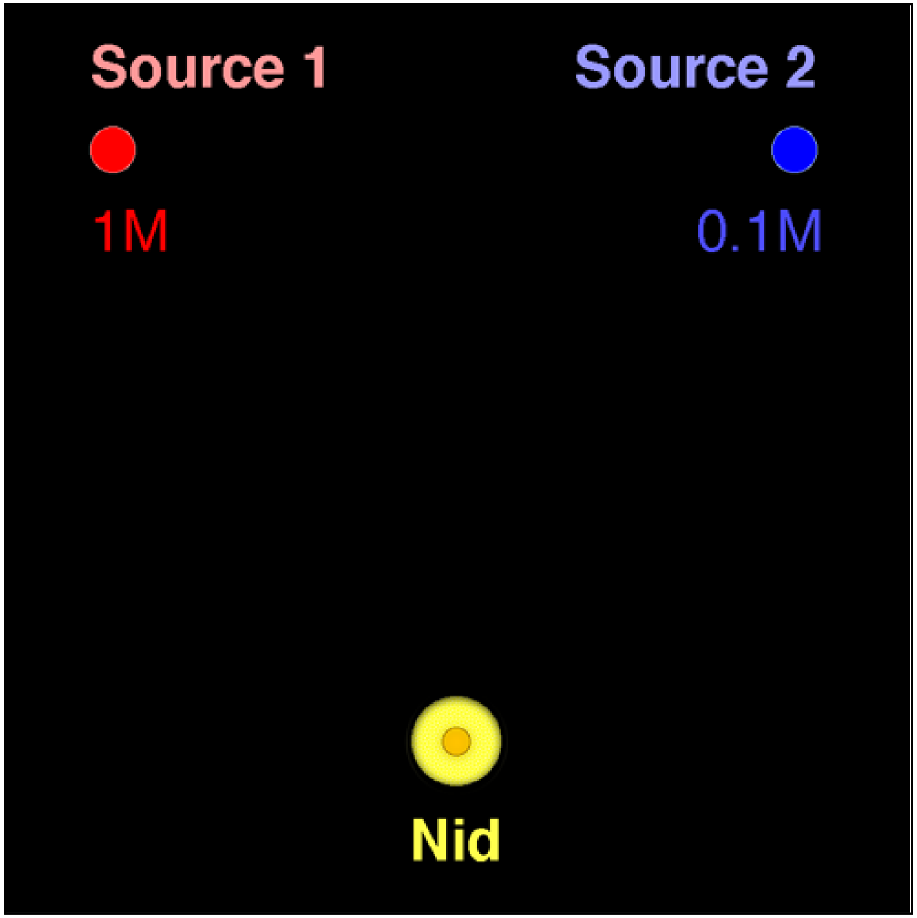

Selecting the Richest Source

Results obtained with *Lasius Niger* ants:

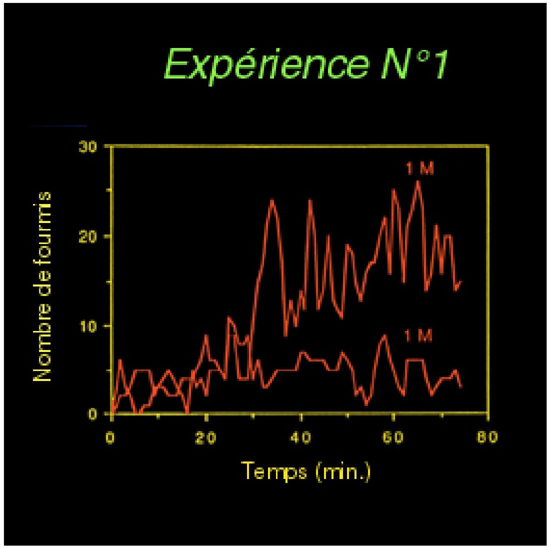

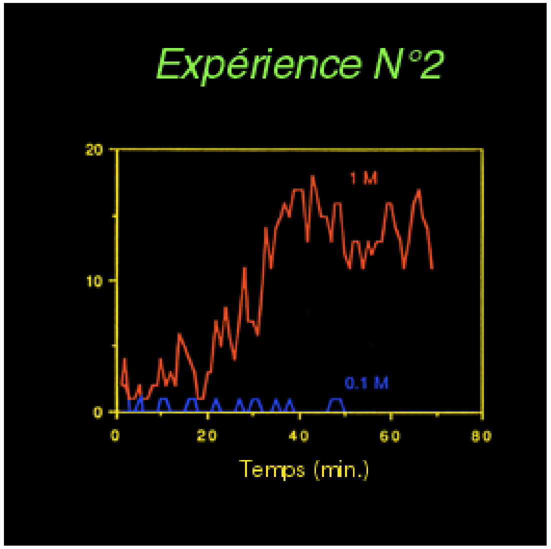

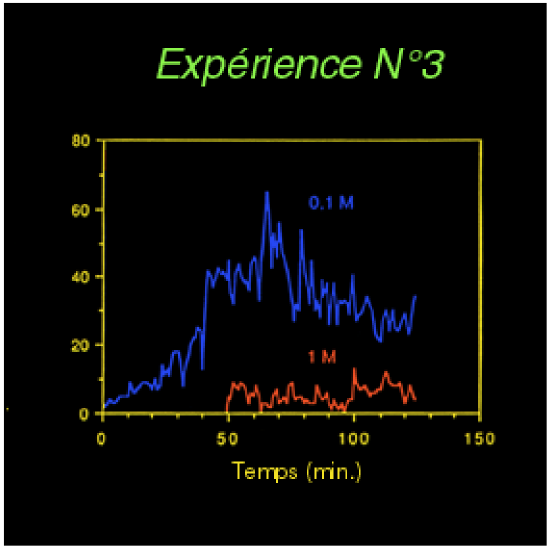

The ants might get stuck within their trail system and therefore the colony exploits primarily the first food source that has been discovered even if this might lead to neglecting a richer source which just appeared at a later time. *Lasius niger:* exclusively uses pheromone-based recruitment mechanisms although has good individual navigation capabilities; probably since nest-sources path not so misaligned, u-turn strategy does not help in this scenario!

Selecting the Richest Source – Scenario 3

Results obtained with *Tetramorium caespitum, Myrmica sabuletti*

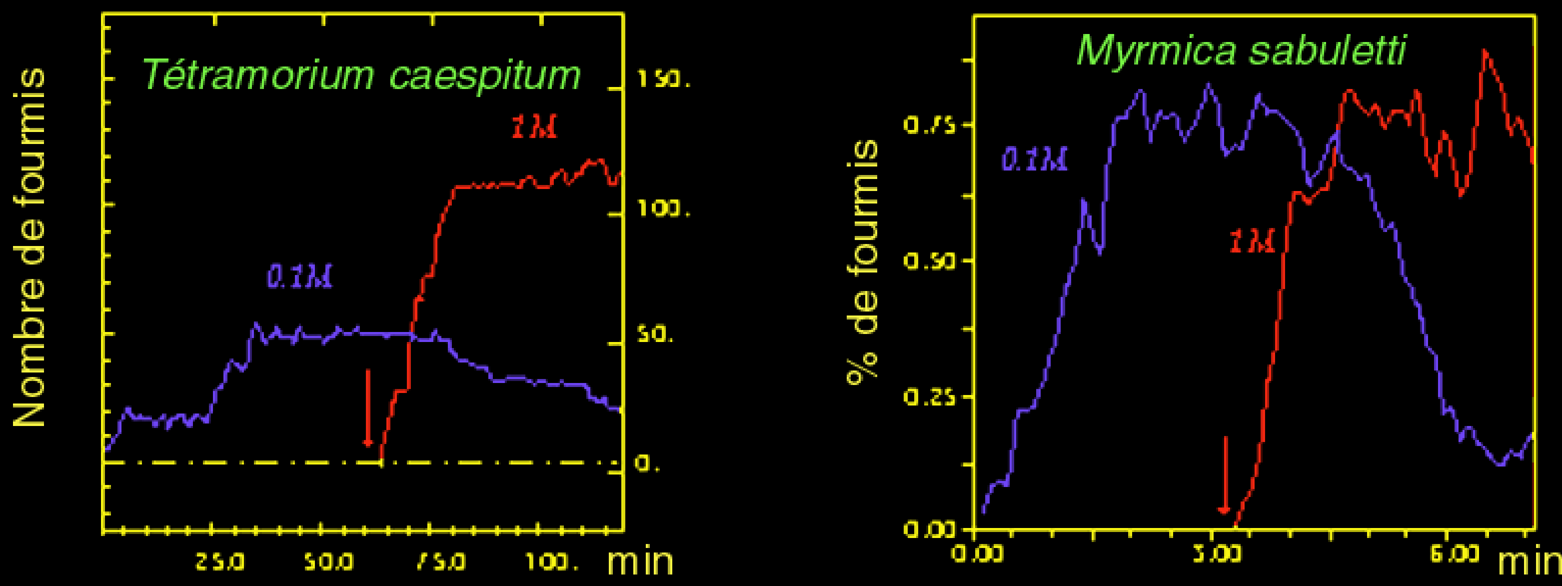

- These two ant species exploit mixed recruitment strategies: mass (trail laying/following) and group (no stigmergy) and do not get stuck in their trail network

Mitchel Resnick, MIT, Media Lab

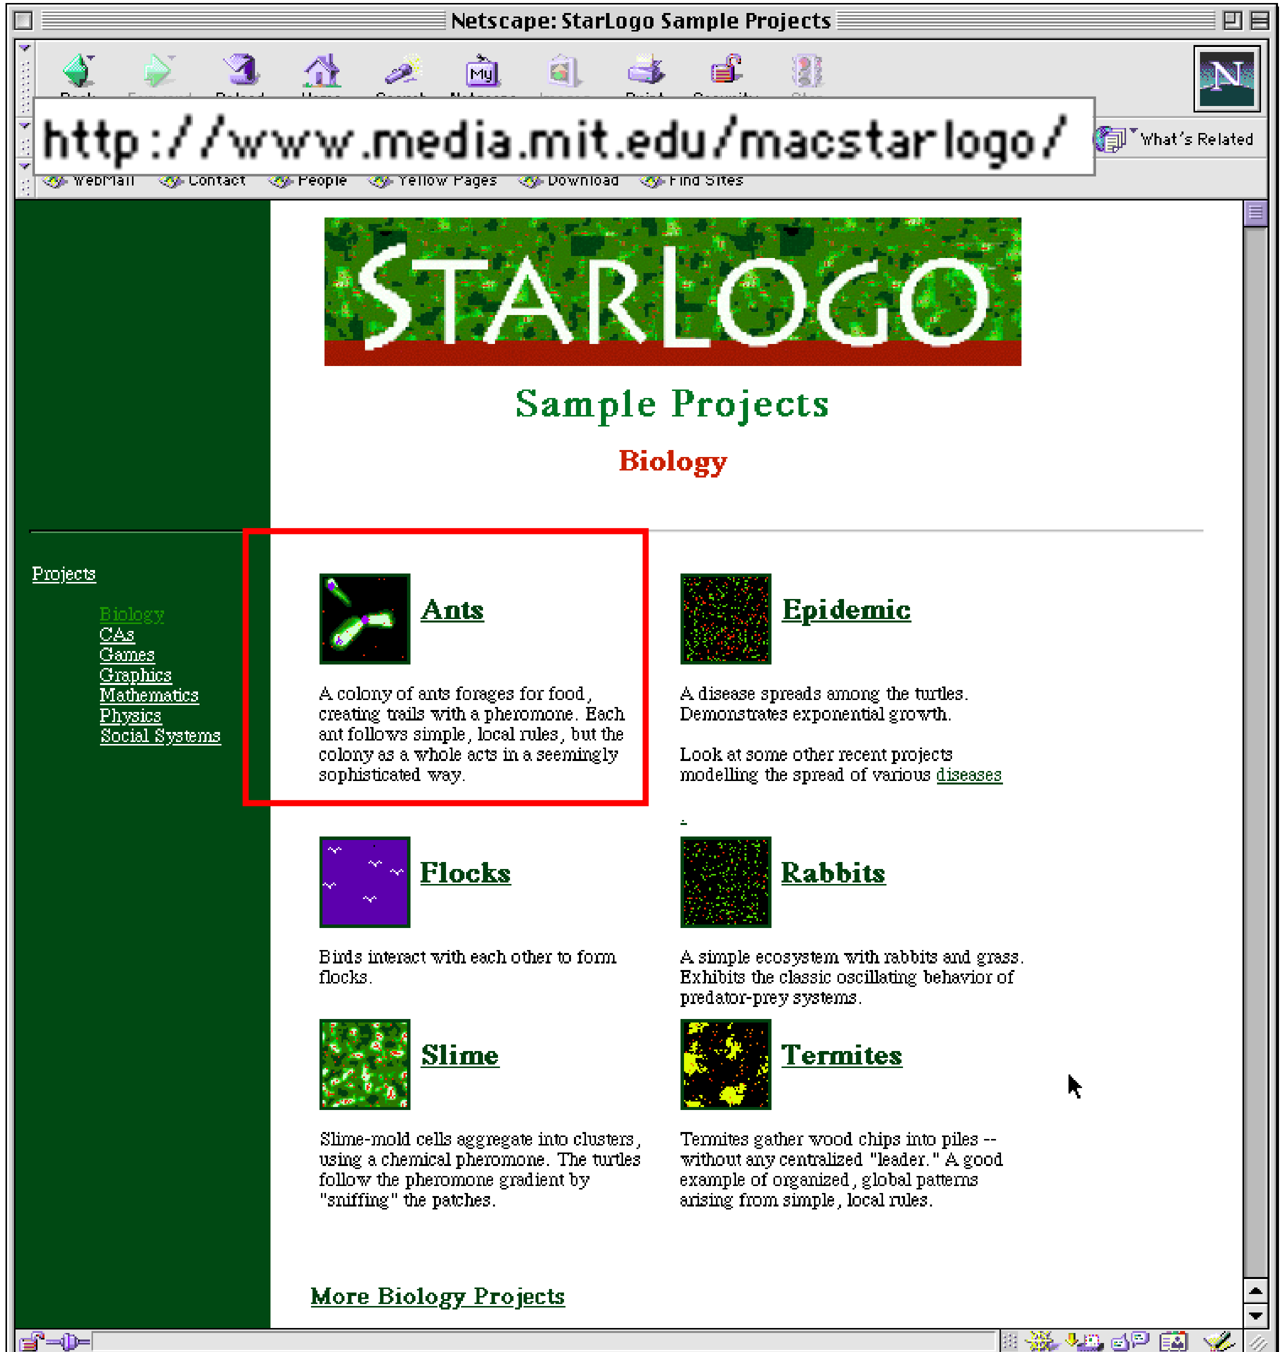

An Example with Three Different Food Sources

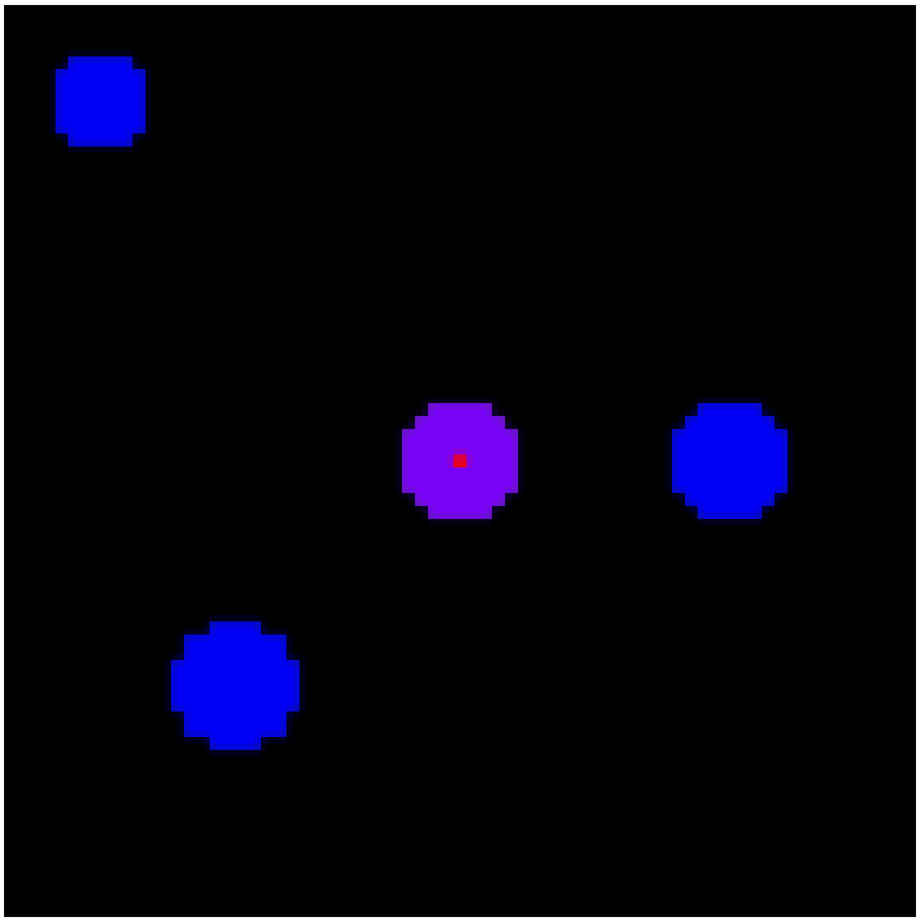

- Different richness
- Different distances from nest
- Obstacle-free environment

Ant Networks

Ant Super-Colonies

The organization of inter-nest traffic in ants

- For most social insects, the fundamental ecological unit is the colony.
- In a number of ant species, groups of workers, larvae, and reproductives can leave the nest and set up a new nest while maintaining close connections with the parent nest.
- The collection of nests, or sub-colonies, forms what is called a super-colony.

Super-colony of *Formica Lugubris* (Switzerland)
Prof. D. Cherix (Uni Lausanne)

The Organisation of Inter-Nest Traffic in Ants

Results for a triangular network (3 nest super-colony)with *Linepithema humile* (Argentine ants)

a,b,c = % of traffic on branch a, b, or c

n = absolute number of passages

[Aron, Deneubourg, Goss, Pasteels, 1991]

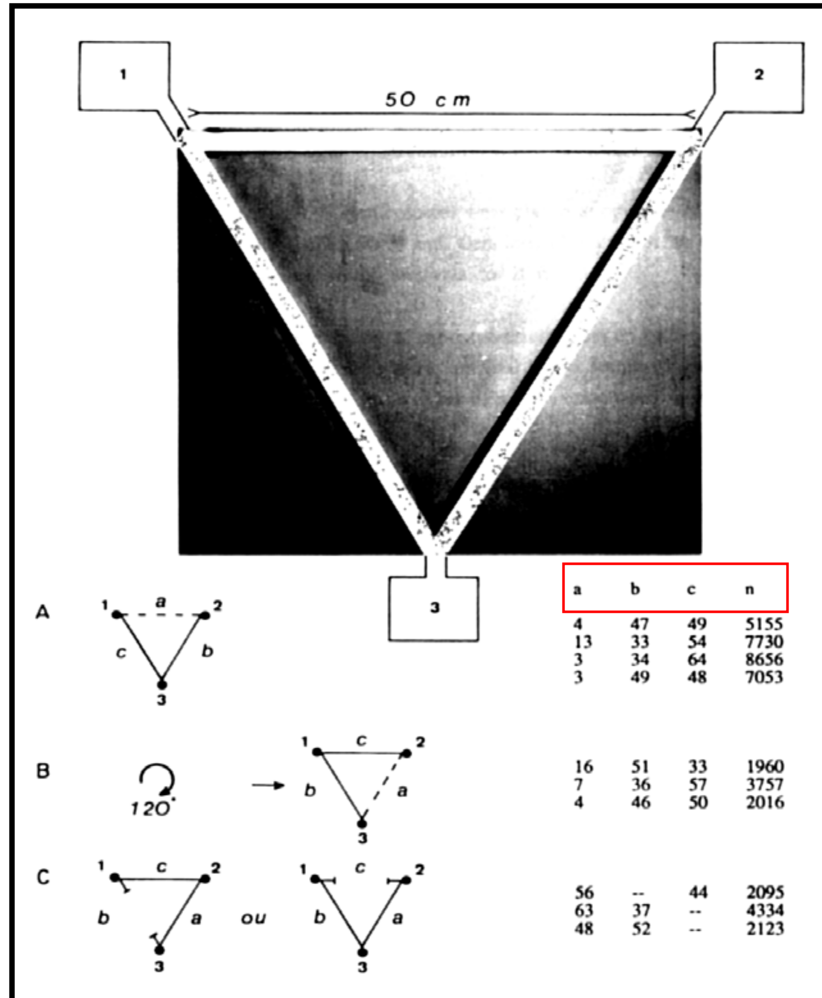

The Organisation of Inter-Nest Traffic in Ants

Results for a quadrangular network (4 nest super-colony) with *Linepithema humile* (Argentine ants)

a,b,c,d = % of traffic on branch a, b, or c

n = absolute number of passages

[Aron, Deneubourg, Goss, Pasteels, 1991]

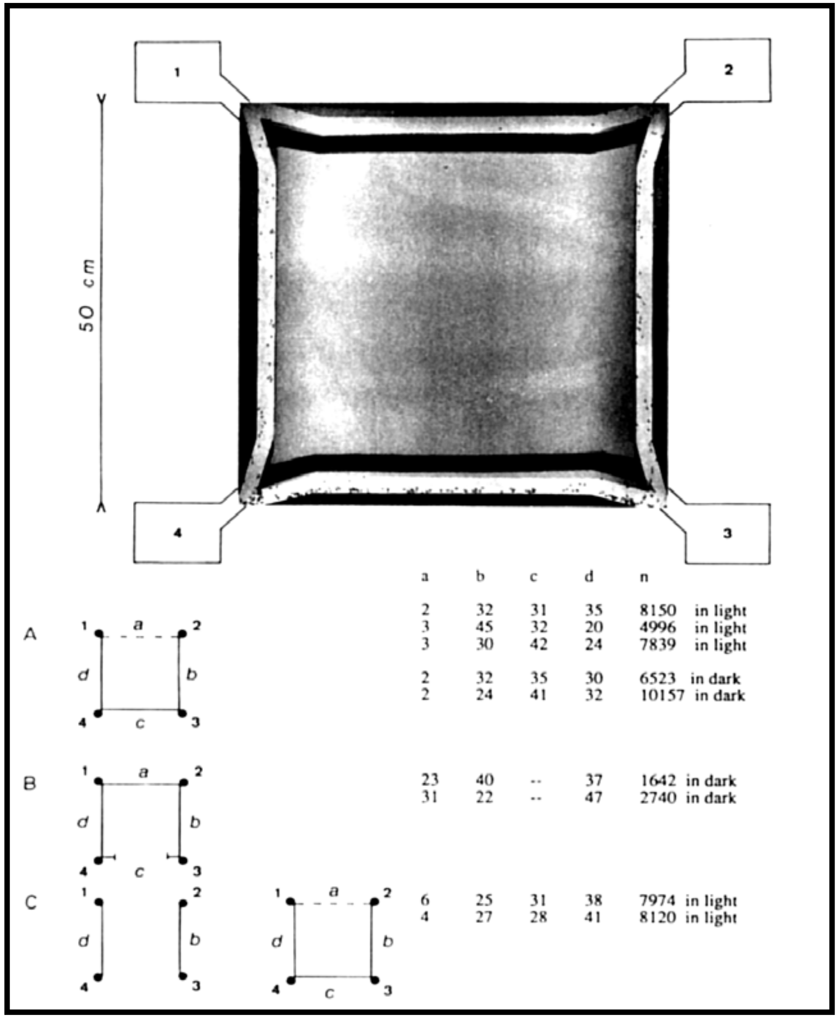

	a	b	c	d	n	
A	2	32	31	35	8150	in light
	3	45	32	20	4996	in light
	3	30	42	24	7839	in light
	2	32	35	30	6523	in dark
	2	24	41	32	10157	in dark
B	23	40	--	37	1642	in dark
	31	22	--	47	2740	in dark
C	6	25	31	38	7974	in light
	4	27	28	41	8120	in light

Ants are Able to Optimize!

- All the nests are connected either directly or indirectly
- Ants are able to find the minimal spanning tree connecting all the nests (probable ecological reasons: cost building and maintaining redundant spanning tree higher + extend predator exposure)
- This is similar to the Traveling Salesman Problem (TSP)
- Can artificial ants solve the TSP?

An Introduction to Multi-Agent Systems based on Ant Trail Laying/Following Mechanisms

Motivation

- Ant Colony Optimization (ACO) algorithms as an example of successful transportation of ideas from natural systems to computational artificial systems (software multi-agent systems)
- ACO algorithms as example of exploitation of swarm intelligence principles as a particular form/instance of distributed intelligence

The Traveling Salesman Problem

The Traveling Salesman Problem

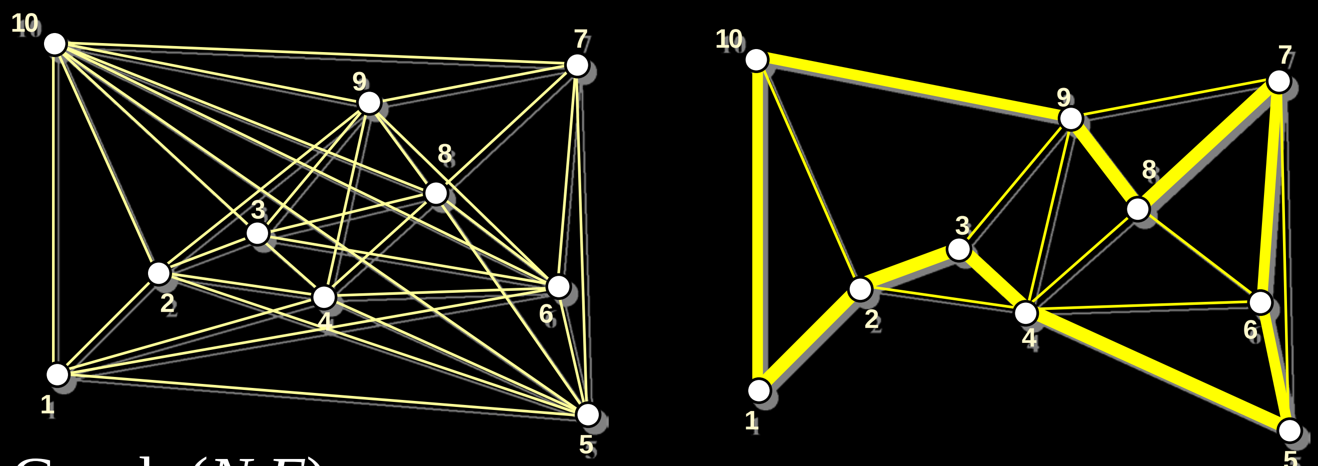

Graph (*N*,*E*)
N : set of cities (nodes)
E : set of connecting roads (links)
d_{ij} : distance between city i and j

Problem: Find the shortest path which allow the salesman to visit once and only once each city in the graph

Difficulty: NP-hard problem; time for computing the shortest route grows in a nonpolinomial way with the number of cities in the network -> metaheuristics/machine-learning class (e.g., ACO, GA) provide near-optimal solutions!

How Hard are NP-Hard Problems?

TSP – Brute force

- A 30 city tour would have to measure the total distance of be 2.65 X 10^{32} different tours. Assuming a trillion additions per second, this would take 252,333,390,232,297 years.
- Adding one more city would cause the time to increase by a factor of 31.

QAP – exact algorithms (e.g. Bixius & Anstreicher 2001)

- around 30+ max instances
- ex. 36 nodes (wiring application): 180h CPU on a 800 MHz Pentium III PC
- Same problem with ACO: 10 s on the same machine

Artificial Ants and the Shortest Path Problem

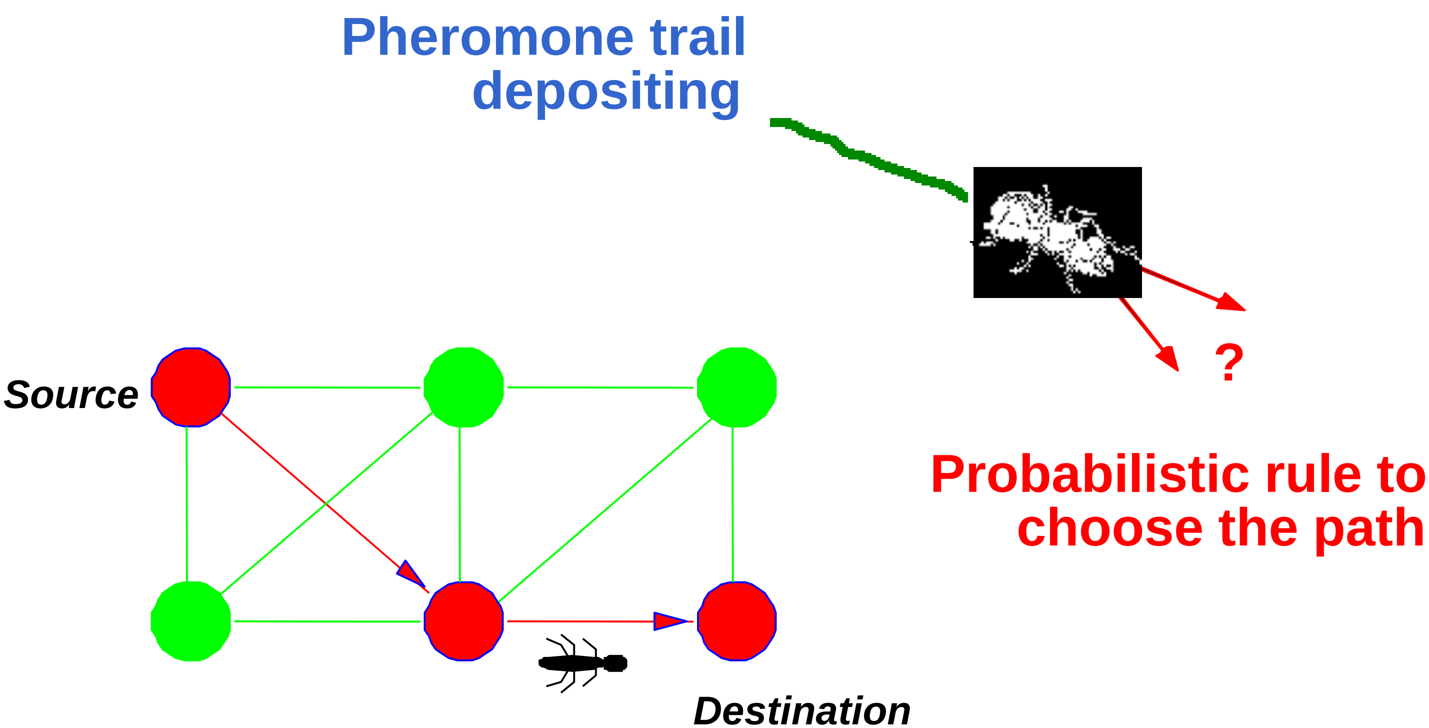

Problem!

The extension of the real ant behavior (forward/backward trail deposit and slow pheromone decay rate) to artificial ants moving on a graph doesn't work:

problem of self-reinforcing loops

Probabilistic solution generation plus pheromone update

-> self-reinforcing loops

Example of possible self-reinforcing loop

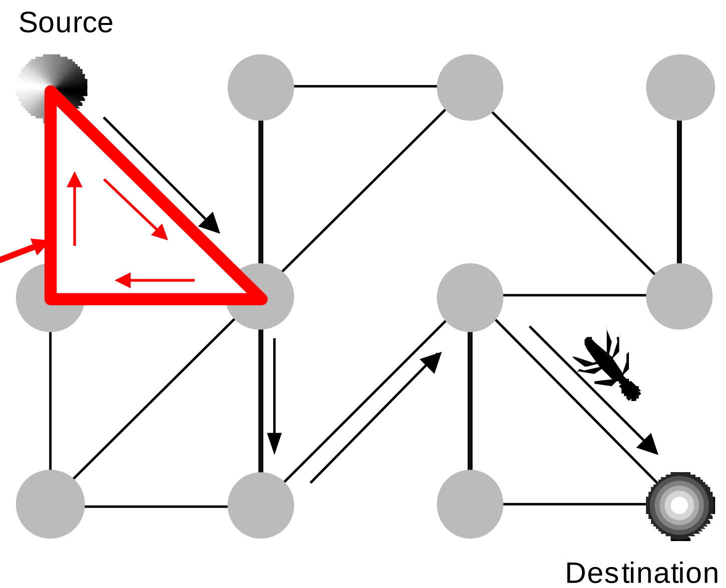

Solution!

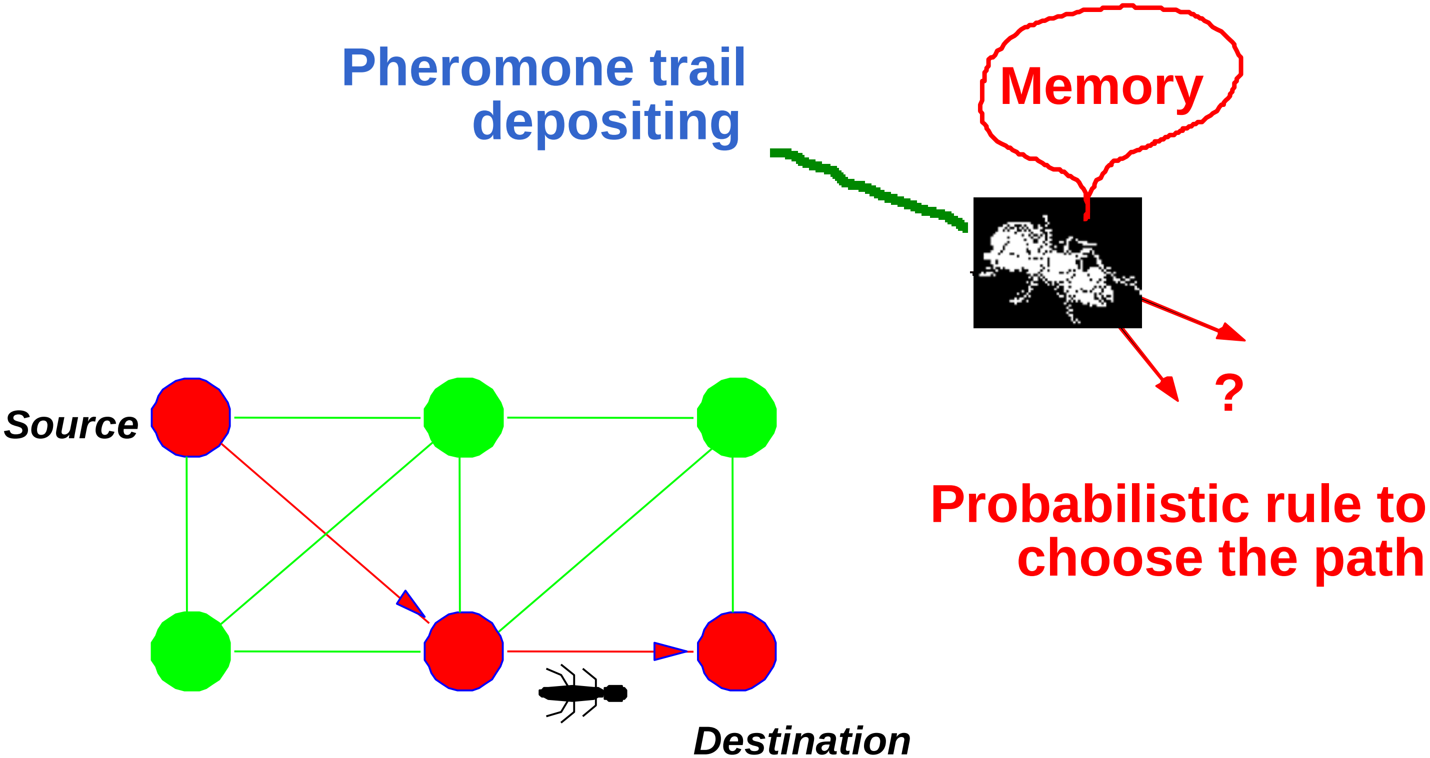
Pheromone trail
depositing
Memory
?
Source
Destination
Probabilistic rule to
choose the path

The First ACO Algorithm: The Ant System (AS)

Design Choices for AS
(Dorigo, Colorni, Maniezzo, 1991)

- Ants are given a memory of visited nodes
- Ants build solutions probabilistically without updating pheromone trails (forwards ants)
- Ants deterministically backward retrace the forward path to update pheromone (backwards ants)
- Ants deposit a quantity of pheromone function of the quality of the solution they generated
- Pheromones evaporates much more quickly than in nature

Assumptions on TSP

- Usual assumption: fully connected graph (i.e. there is a direct route with a given distance from any city in the problem to any other); city list work ok
- Real problem: not fully connected; problem with city list
- Possible solutions:
  - Assume virtual routes so that fully connected; give very bad scores to ants choosing virtual routes (e.g., high but not infinite virtual distance; Dorigo's suggestion)
  - Alternative: break not valid tours asap and either relaunch a new ant or consider less ants for updating pheromones at the next iteration (Martinoli's suggestion); computationally more efficient but risk to lose constructive aspect of trail laying/following; does not work for dead end edges with end criterion being at the start city
  - Graph connectivity: full – dense – sparse; probably different solutions work better as a function of the connectivity degree; interesting problem

AS for TSP - Overview

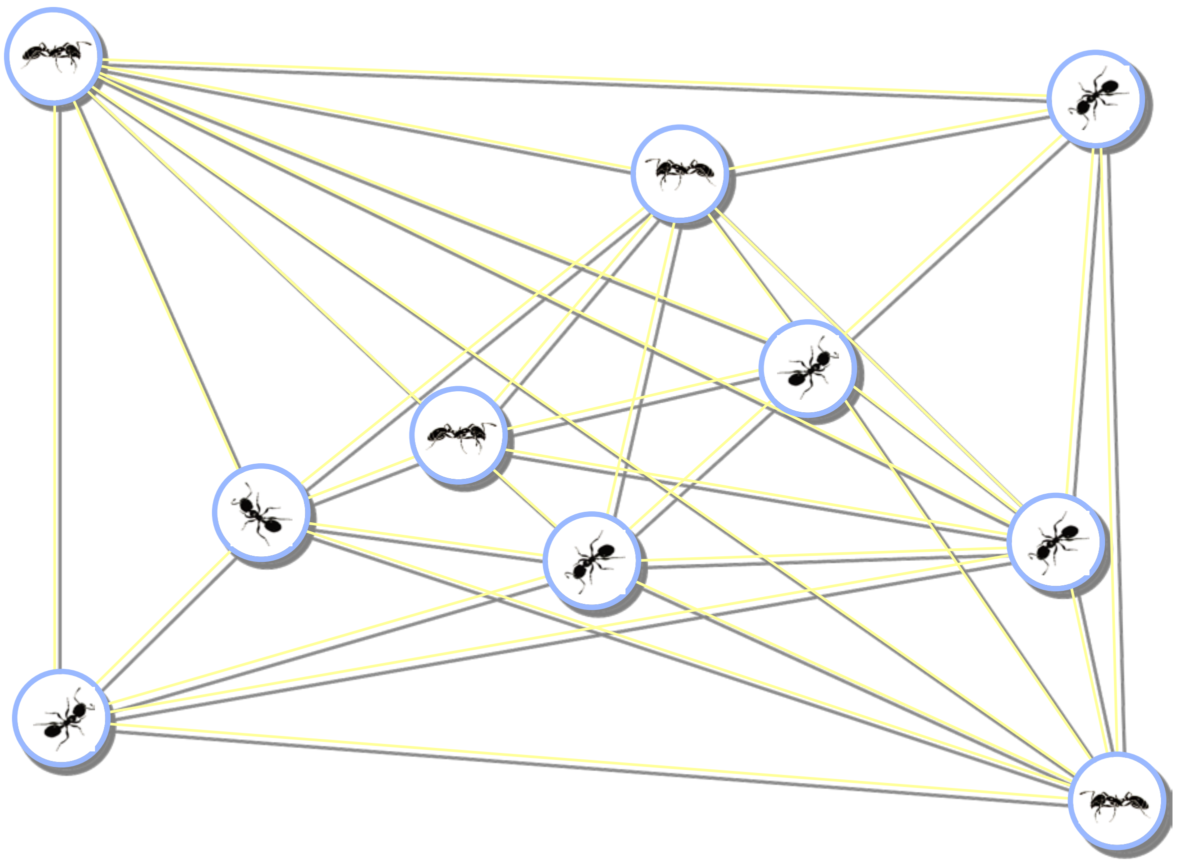

$b_i(t)$, ($i = 1 \ldots n$) : number of ants at the node i at the iteration t

$$m = \sum_{i=1}^{n} b_i(t) = \text{constant: total number of ants}$$

AS for TSP- Individual Ant Behavior

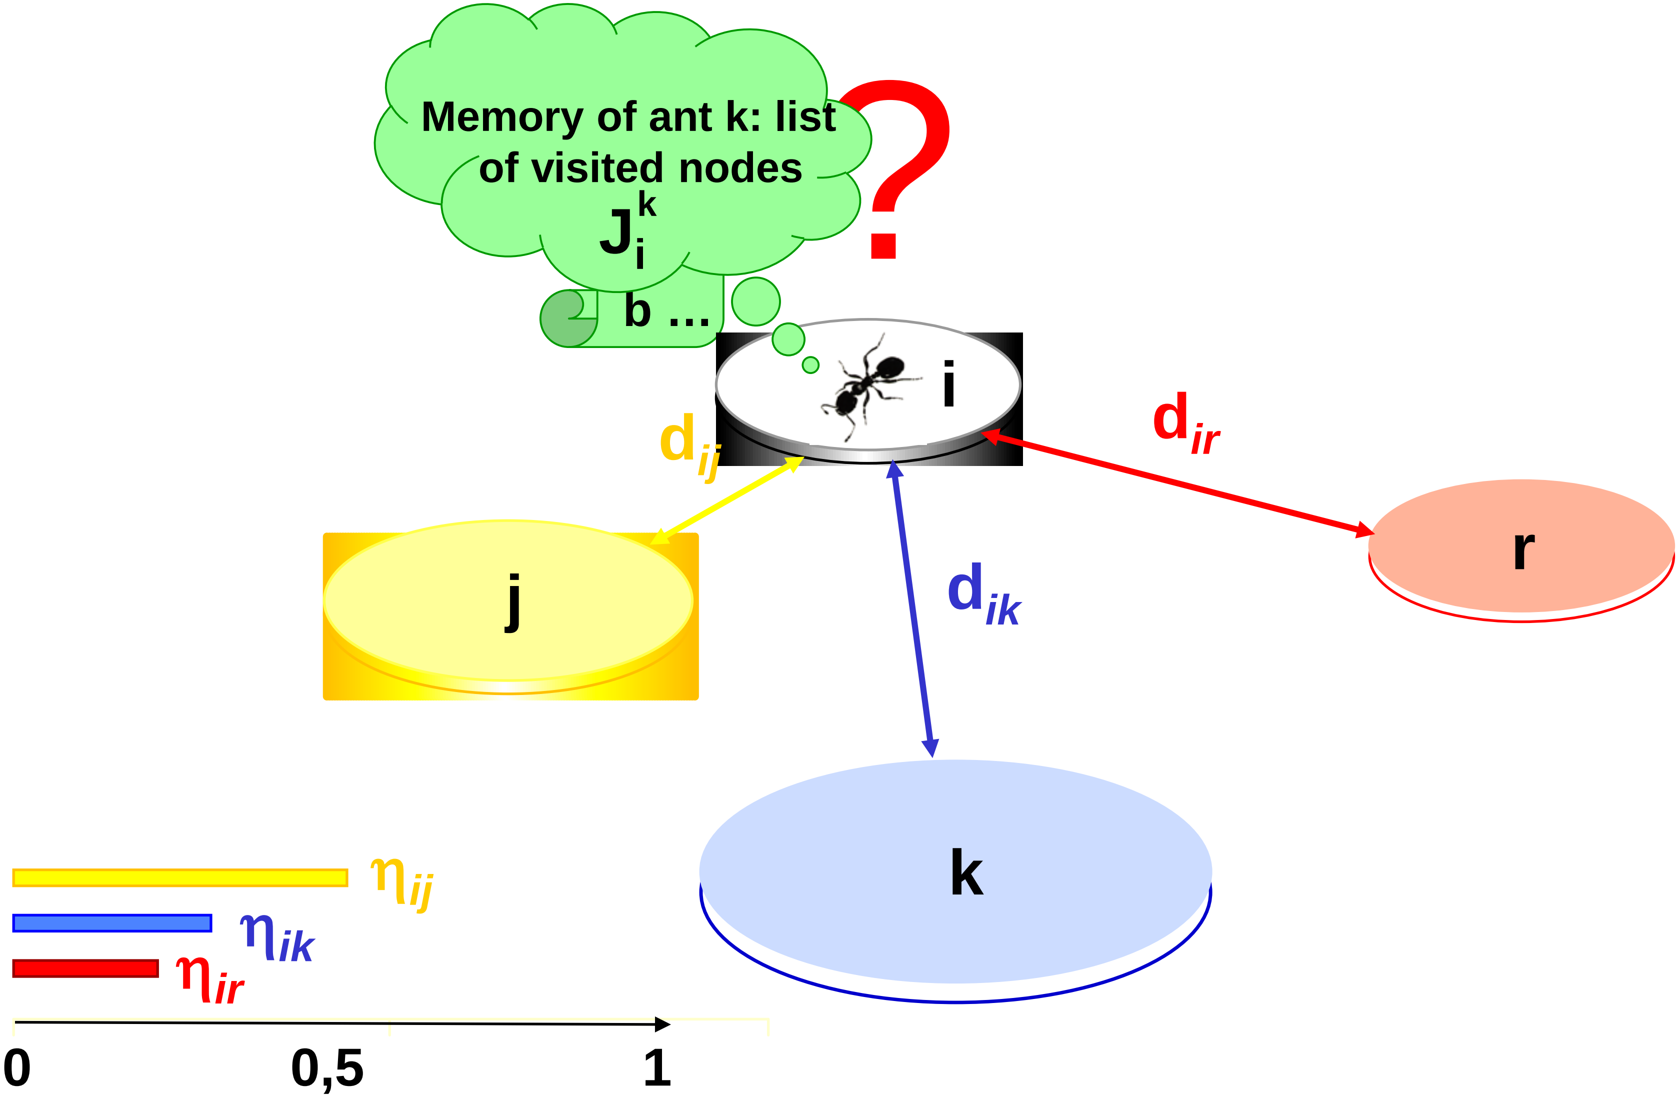

The inverted value of the distance $\eta_{ij} = 1/d_{ij}$ between nodes i and j is called visibility; this information (heuristic desirability) is static, i.e. not changed during the problem solution

AS for TSP- Individual Ant Behavior

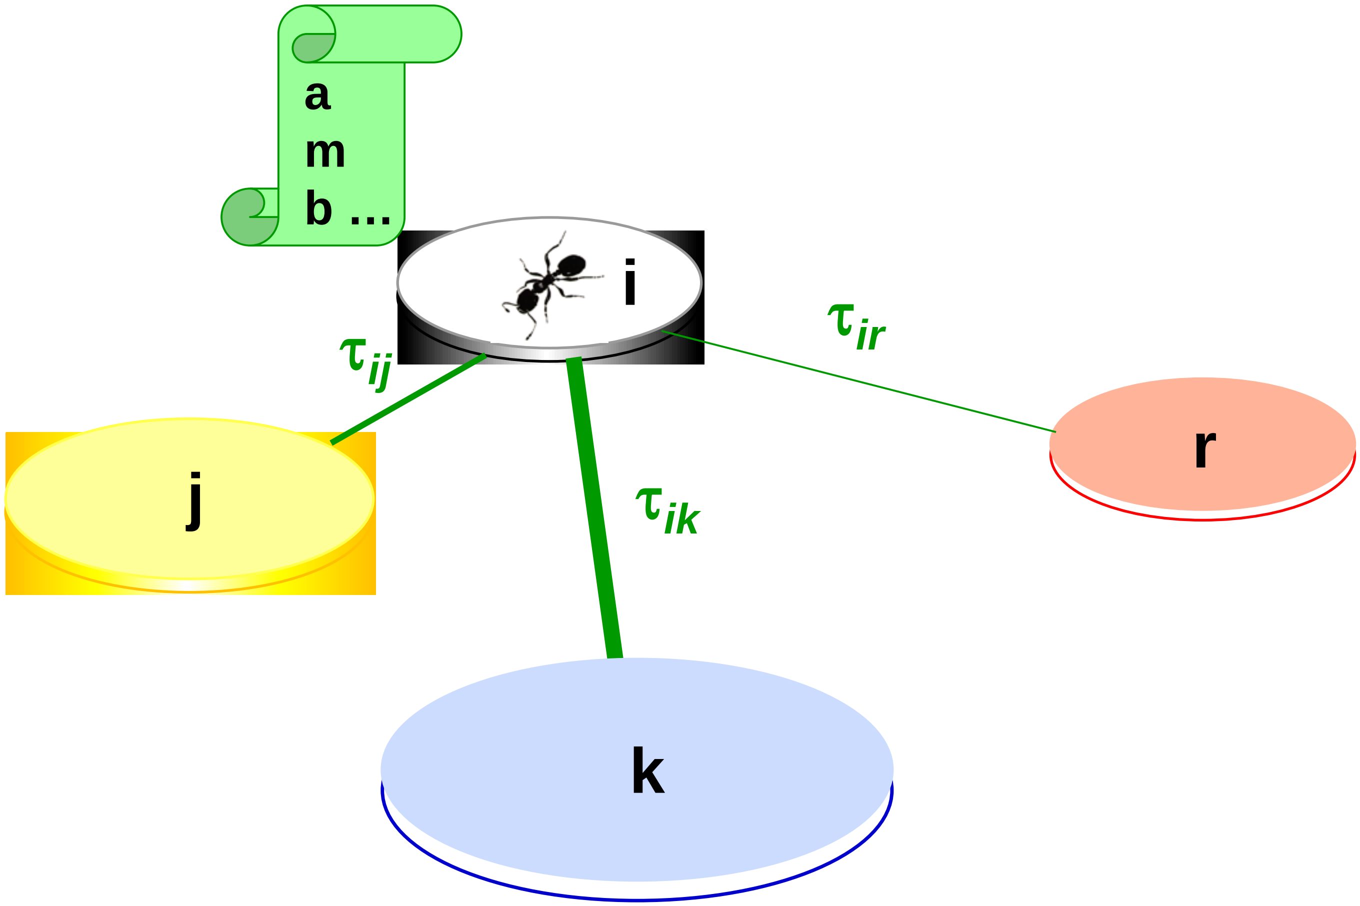

τ_{ij}, quantity of virtual pheromone deposited on the link between the node i and j

AS for TSP - Algorithm

```
Loop \* t = 1*\
    Place one ant on each node \*there are n = |N| nodes \*
    For k := 1 to m \* each ant builds a tour, in this case m=n\*
        For step := 1 to n \* each ant adds a node to its path \*
            Choose the next node to move by applying a
            probabilistic state transition rule
        End-for
    End-for
    Update pheromone trails
Until End_condition \*e.g., t = t_max *\
```

AS for TSP – Transition Rules

J_i^k : list of nodes still to be visited for ant k when it is at node i; starting from an exhaustive list of all the cities in the problem, nodes get scratched during a tour T; at the beginning the list contains all nodes but i; also called tabu list

T : tour, it last n= |N| steps (N = number of nodes in the problem) in which the probabilistic transition rule below is applied

t : iteration index: number of times the whole algorithm is run; $1 \leq t \leq t_{max}$

During a tour T, an ant k at the node i decided to move towards the node j with the following probability (idea: roulette wheel):

$$p_{ij}^k(t) = 0 \quad \text{, if node j have been visited by ant k already because of tabu list}$$

$$p_{ij}^k(t) = \frac{[\tau_{ij}(t)]^{\alpha}\,[\eta_{ij}]^{\beta}}{\sum_{l \in J_i^k} [\tau_{il}(t)]^{\alpha}\,[\eta_{il}]^{\beta}} \quad \text{, if the node have not been visited yet}$$

α : parameter controlling the influence of the virtual pheromone

β : parameter controlling the influence of the local heuristic (visibility)

AS for TSP – Virtual Pheromone Update

At the end of each tour T, each ant k deposits a quantity of virtual pheromone $\Delta\tau_{ij}^{k}$ on the link (i,j); pheromones sum up

$$\Delta\tau_{ij}^{k} = 0 \quad \text{, when (i,j) has not been used during the tour T}$$

$$\Delta\tau_{ij}^{k} = \frac{Q}{L^{k}(t)} \quad \text{, when (i,j) has been used during the tour T}$$

$L^{k}(t)$ = length of the tour T done by ant k at iteration t
Q = parameter (adjusted by heuristic, not sensitive)

Note: the longer the tour, the lower is the quality of the solution, the smaller the quantity of pheromone dropped

AS for TSP – Default Virtual Pheromone Update

$$\tau_{ij}(t+1) \leftarrow (1-\rho)\tau_{ij}(t) + \Delta\tau_{ij}(t)$$

with $\Delta\tau_{ij}(t) = \sum_{k=1}^{m} \Delta\tau_{ij}^{k}$

ρ = evaporation coefficient

At iteration t = 0 each link is initialized with a small homogenous pheromone quantity τ_0

AS for TSP – Virtual Pheromone Update with Elitism (EAS)

$$\tau_{ij}(t+1) \leftarrow (1-\rho)\tau_{ij}(t) + \Delta\tau_{ij}(t) + e\Delta\tau_{ij}^{e}(t)$$

with $\Delta\tau_{ij}(t) = \sum_{k=1}^{m} \Delta\tau_{ij}^{k}$

$\Delta\tau_{ij}^{e}(t) = Q / L^{+}$ if (i,j) belongs to the best tour T^{+} out of the m tours generated by ants at a given iteration

$\Delta\tau_{ij}^{e}(t) = 0$ otherwise

e = parameter (adjusted by heuristic, not sensitive)

Note: idea, best tours get extra reinforcement

AS for TSP – Evolution of the Best Tour Length

Example: 30 nodes problem

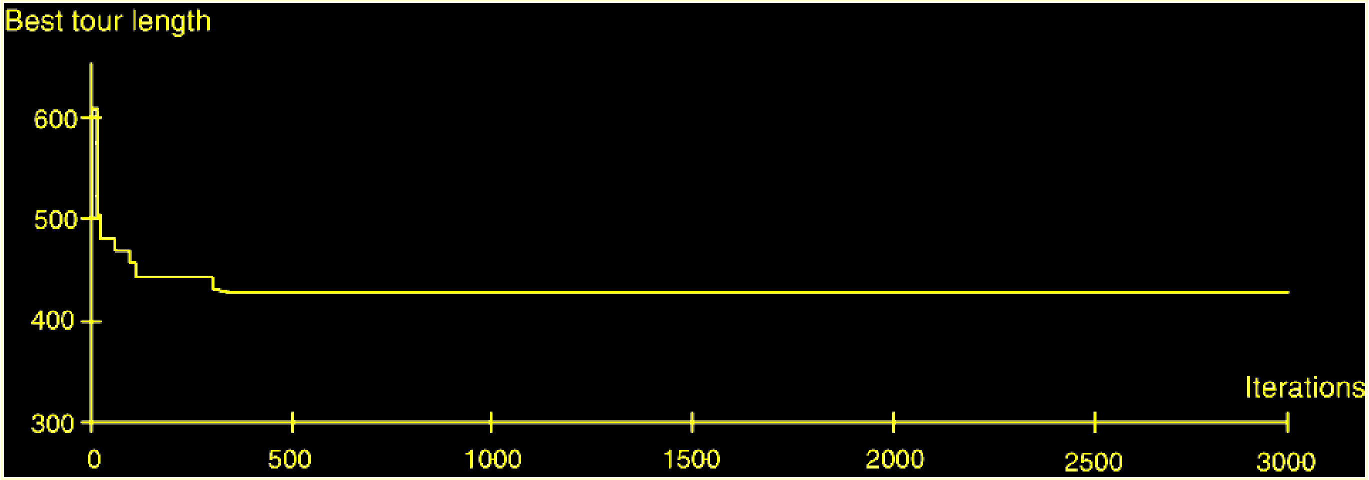

AS for TSP – Results 50 cities

Example of solution found on Eil50 problem

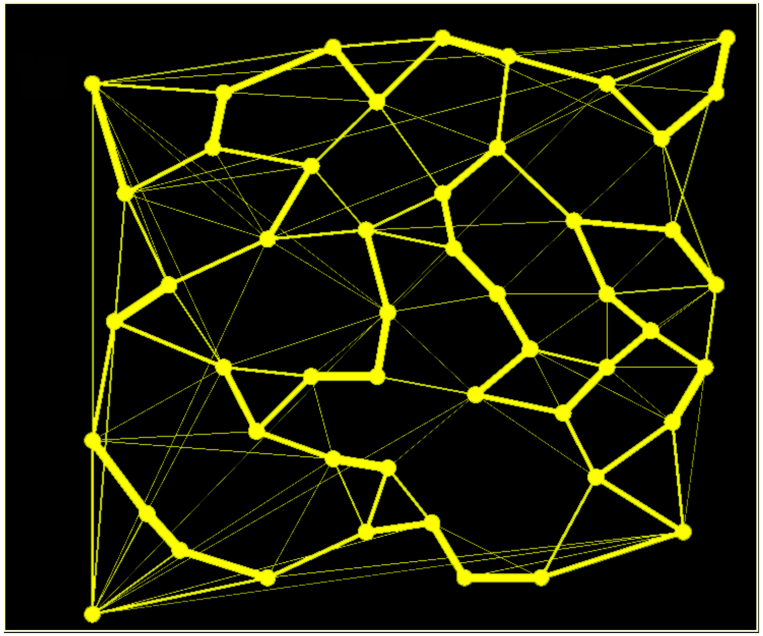

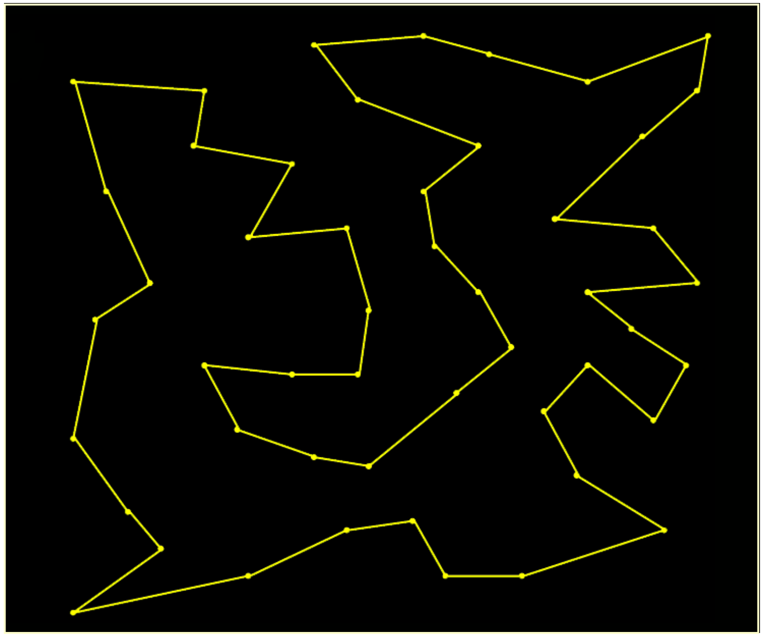

AS for TSP – Performance as a Function of the Problem Dimension

Network	n (dimension)	best solution	Mean number of iterations for to the near-optimal solution	Simulation time (seconds)
4 X 4	16	160	5,6	8
5 X 5	25	254,1	13,6	75
6 X 6	36	360	60	1020
7 X 7	49	494,1	320	13440
8 X 8	64	640	970	97000

Summary of AS

- **Ants** are launched at each iteration from each node to explore the network
- **Ants** build their paths probabilistically with a probability function of:
  - (i) **artificial pheromone values**, and
  - (ii) **heuristic values** (in TSP: city visibility)
- **Ants** memorize visited nodes
- Once reached their destination nodes (in TSP the last city on their list) **ants** retrace their paths backwards, and update the pheromone trails

Conclusion

Take Home Messages

1. Differences between artificial and natural SI
2. Differences between computational and physical SI
3. Key mechanisms for natural SI: self-organization and stigmergy
4. Self-organization ingredients: positive feedback, negative feedback, randomness, multiple interactions
5. SI-based systems exploit careful balance between exploration and exploitation
6. Microscopic models help understanding SI-based systems
7. Ants exploit trail laying/following mechanisms and other strategies for foraging

Take Home Messages

8. Ants are able to generate efficient inter-nest networks
9. Trail laying/following mechanisms can be expanded with other properties of the agent easily implementable in software (e.g., memory, modulation of the pheromone quantity, etc.)
10. Ant System has been the first metaheuristic taking advantage of the ant inspiration
11. The first NP hard problem it has been applied was the Traveling Salesman Problem

Additional Literature – Week 1

Books

- Kelly K., "Out of Control: The New Biology of Machines, Social Systems and the Economic World", Perseus Press, 1995.
- Ormerod P., "Butterfly Economics: A New General Theory of Social and Economic Behavior", Pantheon Books, 2001.
- Crichton M., " Prey: A Novel", HarperCollins, 2002.
- Bonabeau E. and Theraulaz G., "Intelligence Collective". Hermès, 1994 (in French).
- Camazine S., Deneubourg J.-L., Franks N. R., Sneyd J., Theraulaz G., Bonabeau E., "Self-Organization in Biological Systems". Princeton Studies in Complexity, Princeton University Press, 2001.
- Hölldobler B. and Wilson E. O., "Journey to the Ants: A story of Scientific Exploration", The Harvard University Press, Cambridge, MA, 1994.

Additional Literature – Week 1

Papers

- Beni G. and Wang J., “Swarm Intelligence”. *Proc. of the Seventh Annual Meeting of the Robotics Society of Japan*, Tokyo, Japan, 1989, pp. 425-428.
- Bonabeau E. and Theraulaz G. "Swarm Smarts". *Scientific American*, March 2000, pp. 72-79.
- Peter Miller “Swarm Theory”, *National Geographic*, July 2007, pp. 126-147.
- Aron S., Beckers R., Deneubourg J.-L., and Pasteels J.-M., “Memory and Chemical Communication in the Orientation of two Mass-Recruiting Ant Species”. *Insect Society*, 1993, Vol. 40, pp. 369-380.
- Deneubourg J.-L., Aron S., Goss S., and Pasteels J. M., “The Self-Organizing Exploratory Pattern of the Argentine Ant”. *J. of Insect Behavior*, 1990, Vol. 3, pp. 159-168.
- Beckers R., Deneubourg J.-L., and Goss S., “Trails and U-Turns in the Selection of a Path by the Ant *Lasius niger*”. *J. Theor. Biol.*, Vol. 159, pp. 397-415, 1992.
- Jackson D. E., Holcombe M., and Ratnieks F. L. W., “Trail geometry gives polarity to ant foraging networks”. *Nature*, No. 432, pp. 907-909, 2004.
- Dorigo M., Maniezzo V., and Colorni A., “The Ant System: Optimization by a Colony of Cooperating Agents”. *IEEE Trans. Systems, Man, and Cybernetics, Part B*, Vol. 26, pp. 29-41, 1996.