Stochastic optimization algorithms
Lecture 13, 20181003

Ant colony optimization (II)

Mattias Wahde, PhD, Professor, Chalmers University of Technology
e-mail: mattias.wahde@chalmers.se, http://www.me.chalmers.se/~mwahde

Exam information

- Sign up for the exam! Last day (to sign up): 20181011
- Note: It is ***mandatory*** to sign up for the exam!
- Exam date: 20181031, 14.00-18.00, M
- More information about the exam, as well as some practice exams, will follow next week (Tuesday).

FAQ

- Check the FAQ (on the course web page) regularly!
- For example, some important clarifications regarding HP2 have already been added, and more might be added later.
- When correcting the home problems, we will assume that you have read the FAQ carefully.

Node selection in ACO

- (See also the previous lecture, 20181002)
- The starting node is selected randomly (for each ant) in order to maximize the chances of finding a good path, by exploring the search space as thoroughly as possible.
- In general (after selecting the start node), the move from node j to no i is carried out probabilistically, in proportion to $p(e_{ij}|S)$.

Node selection in ACO (example)

- Three available links, with probabilities $p(e_{ij}|S)$ equal to 0.7, 0.2 and 0.1, respectively.
- Method for selecting the edge to traverse:
  - Draw a random number r.
  - If r $\leq$ 0.7, choose the first edge (right).
  - If 0.7 $<$ r $\leq$ 0.9, choose the second edge (middle).
  - Otherwise (r $>$ 0.9), choose the third edge (left).

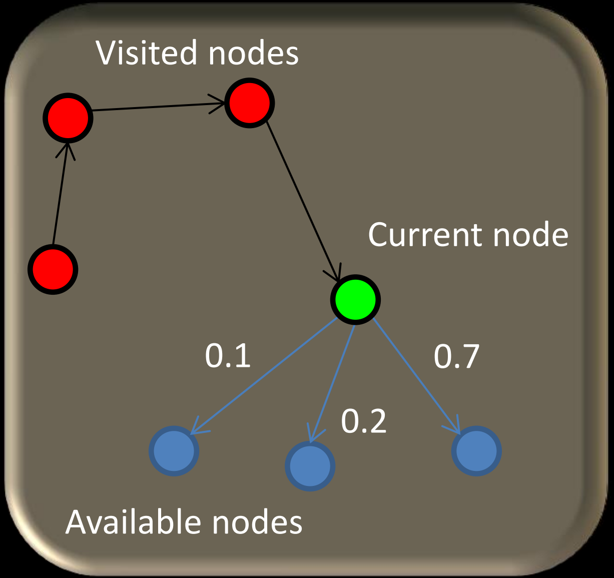

Today's learning goals

- After this lecture you should be able to
  - Describe and implement the MMAS algorithm
  - Prove that pheromone levels are limited in MMAS
  - Prove that MMAS converges
  - Describe some aspects of cooperative robotics

Max-min ant system (MMAS)

- Similar to AS, but there are three major differences
- In MMAS (Algorithm 4.2)
  – Only the ant generating the best soluton is allowed to deposit pheromone,
    - “best” = best-in-current-iteration *or* best-so-far. Both options possible.
  – Pheromone limits ($\tau_{\min}$ and $\tau_{\max}$) are imposed to avoid stagnation,
  – Pheromone levels initialized as $\tau_{\max}$.

Note: Misprint after Eq. (4.7), should say $\Delta\tau_{ij}{}^{[b]}$, not $\tau_{ij}{}^{[b]}$.

Mattias Wahde, PhD, Professor, Chalmers University of Technology
e-mail: mattias.wahde@chalmers.se, http://www.me.chalmers.se/~mwahde

Today's learning goals

- After this lecture you should be able to
  - Describe and implement the MMAS algorithm
  - Prove that pheromone levels are limited in MMAS
  - Prove that MMAS converges
  - Describe some aspects of cooperative robotics

Mattias Wahde, PhD, Professor, Chalmers University of Technology
e-mail: mattias.wahde@chalmers.se, http://www.me.chalmers.se/~mwahde

Pheromone limits in MMAS

- In MMAS, the pheromone level on any edge has an upper bound equal to f^*/ρ where f^* is the value of the objective function for the best solution (e.g. $1/D^*$ in the case of TSP).
- Proof (presented on the blackboard), see pp. 183-184.
- In MMAS, $\tau_{\max}$ is initially set as $1/D^{nn}$, and is then updated (along with $\tau_{\min}$) during a run, whenever a new best path is found.
- $\tau_{\min}$ is set (empirically) as $\frac{\tau_{\max}\left(1-\sqrt[n]{0.05}\right)}{\left(\frac{n}{2}-1\right)\sqrt[n]{0.05}}$.

Mattias Wahde, PhD, Professor, Chalmers University of Technology
e-mail: mattias.wahde@chalmers.se, http://www.me.chalmers.se/~mwahde

Today's learning goals

- After this lecture you should be able to
  - Describe and implement the MMAS algorithm
  - Prove that pheromone levels are limited (from above) in MMAS
  - Prove that MMAS converges
  - Describe some aspects of cooperative robotics

Convergence proof for MMAS

- Let $p(k)$ = the probability of encountering the best solution at least once in the first k iterations.
- One can then prove (for MMAS) that

$$\lim_{k \to \infty} p(k) = 1$$

- The proof (see p. 184) is simple, and rests on the fact that pheromone levels are bounded in MMAS, so that all solutions have non-zero probability.

Today's learning goals

- After this lecture you should be able to
  - Describe and implement the MMAS algorithm
  - Prove that pheromone levels are limited in MMAS
  - Prove that MMAS converges
  - Describe some aspects of cooperative robotics

Comparison with AS

- For TSP, MMAS often achieves better *average* results than AS, but with larger spread (standard deviation).
- When running MMAS, one can improve performance by …
  - … restarting the run, re-initializing the pheromone limits (minimum and maximum) based on the most recent available estimate of $\tau_{\max}$.
  - … alternating the definition of the best solution from time to time (sometimes using best-so-far, sometimes best-in-current-iteration).
- For HP2.1c, use standard AS, not MMAS.

Applications

- TSP-related problems – vehicle routing, network routing, component placement.
  - NOTE: ACO good at handling dynamic problems, i.e. problems in which conditions change *during* a run.
- Scheduling (see pp. 112-114), read by yourselves.

Mattias Wahde, PhD, Professor, Chalmers University of Technology
e-mail: mattias.wahde@chalmers.se, http://www.me.chalmers.se/~mwahde

Vehicle routing problem (VRP)

- Static:
  - Problem: Use a fleet of N vehicles, each with capacity Q_i to visit a number (n) of customers, deliver a quantity q_k of goods to each customer, taking time t_k.
  - Goal: Find a tour with the shortest possible total time.
- Dynamic (DVRP):
  - New orders arrive continuously
  - The aim is to serve all known orders at any given time.
- This problem has been approached by many authors, using various versions of ACO.

Cooperative robotics

- Not really an *application* of ACO *per se*, but strongly related.
- Goal: To make a set of identical (or at least similar) robots carry out complex tasks (such as search-and-rescue), beyond the capability of a single robot.
- Several important subfields of robotics have emerged, e.g. *rescue robotics*.

Cooperative transport by robots

- Cooperative box-pushing (Kube and Zhang)
- The box emitted (visible) light, detectable by the robots.
- Robots had five elementary behaviors:
  - Find (for locating the box)
  - Slow (for slowing down, to avoid rear-end collisions)
  - Follow (for following another robot)
  - Avoid (for avoiding collisions)
  - Goal (for directing a robot towards the box).

Cooperative transport by robots

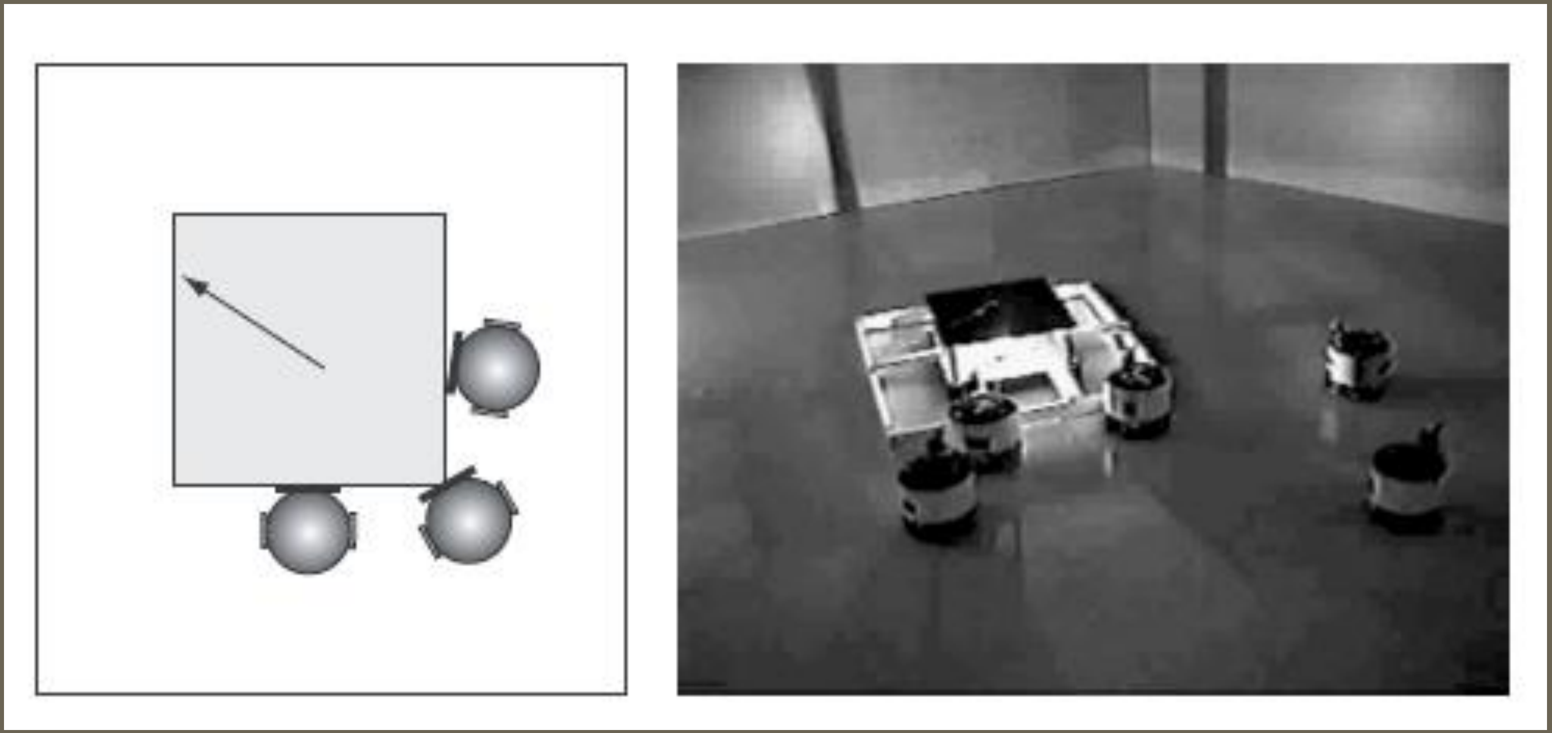

The Swarm-bot project

- Led by M. Dorigo.
- Goal: To develop a set of robots that together could carry out complex tasks (such as search-and-rescue) beyond the capability of a single robot.
- Several ant-like complex behaviors were generated, e.g. dynamic bridge-building and cooperative object transportation.

The Swarm-bot project

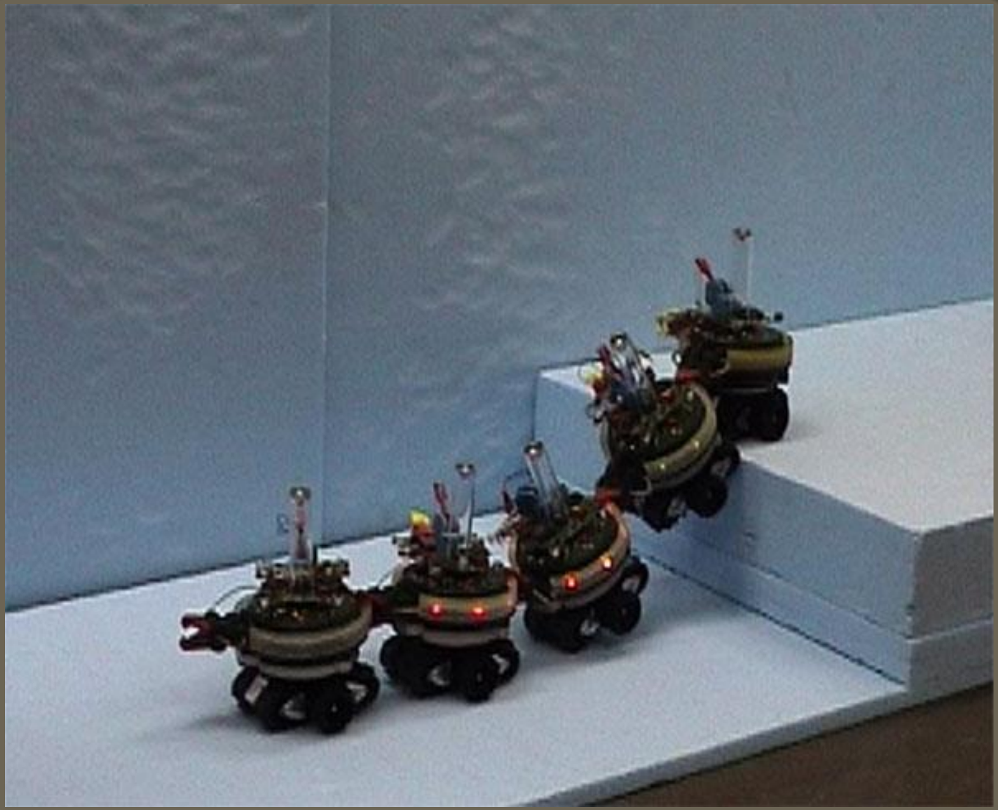

Harvard's kilobot swarm

- A recent (2014) example of swarm robotics.
- The swarm contained 1024 simple, three-legged robots.
- Using only local interactions, the robots can (collectively) form various shapes, while avoiding deadlocks.

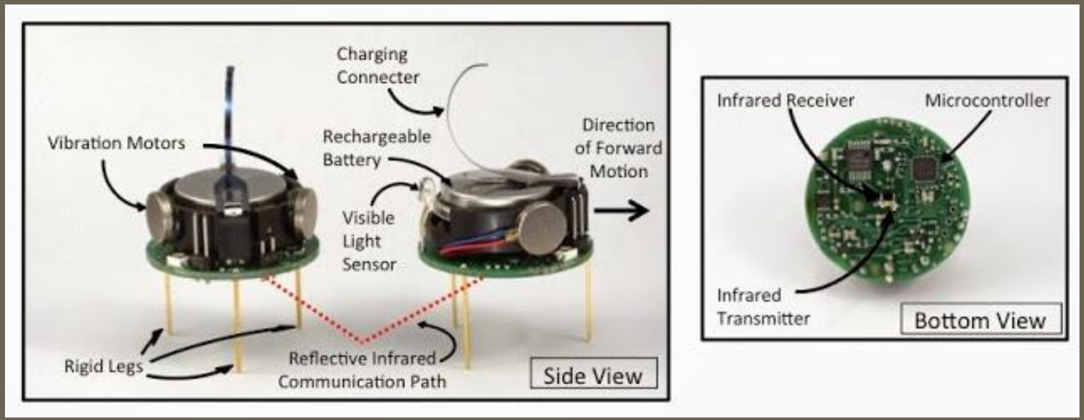

Harvard’s kilobot swarm

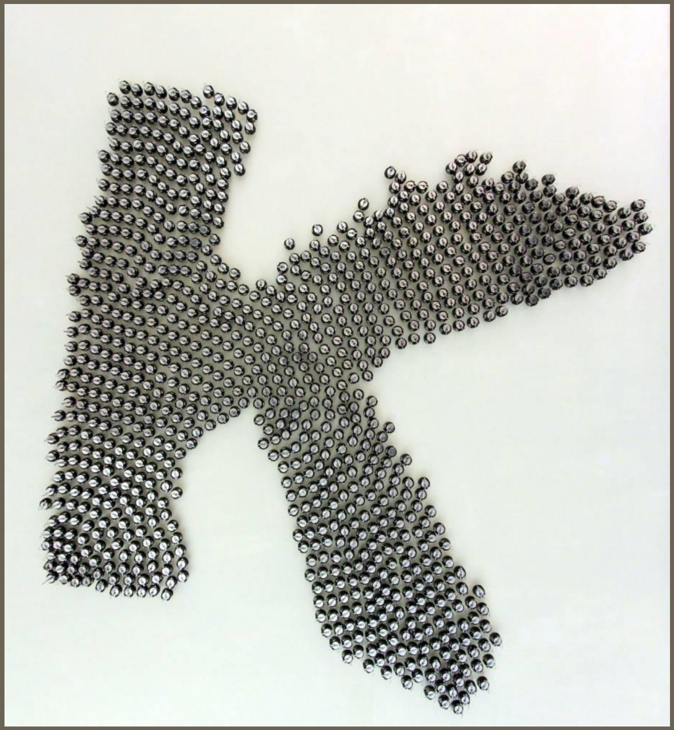

Today's learning goals

- After this lecture you should be able to
  - Describe and implement the MMAS algorithm ✓
  - Prove that pheromone levels are limited in MMAS ✓
  - Prove that MMAS converges ✓
  - Describe some aspects of cooperative robotics ✓