Slide 0

Ant Colony Optimization

Reference book: "Ant Colony Optimization" by Marco Dorigo and Thomas Stutzle (The MIT Press)

Slide 1

Background

- social insect behavior →swarm intelligence
  - large assemblies of simple units solve complex tasks
  - no central control
  - self-organized behavior
  - emergent behavior
  - ☛ swarm intelligence
  - social tasks
    - foraging, defense & attack, nest building etc.
  - some of these tasks can be viewed as "natural optimization problems"
    - e.g., foraging as optimal resource usage (food vs. energy waste)

Background (Cont'd)

Slide 2

- swarm intelligence
  - Collective system capable of accomplishing difficult tasks in dynamic and varied environments without any external guidance or control and with no central coordination
  - Achieving a collective performance which could not normally be achieved by an individual acting alone
  - Constituting a natural model particularly suited to distributed problem solving



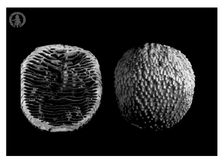


Background (Cont'd)

Slide 3

- Natural behavior of ant

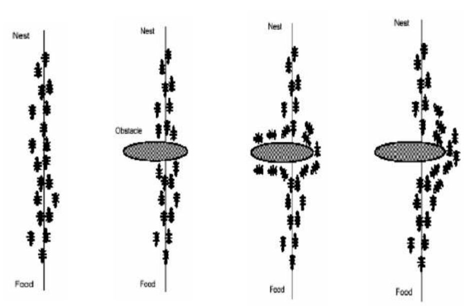

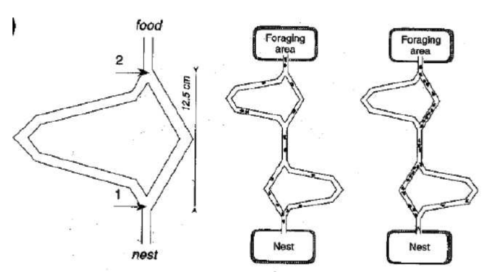

Slide 4

From real to artificial ants

- ant colonies
  - in spite of the simplicity of their individuals,
  - present highly structured social organization
  - can accomplish complex tasks that far exceed the individual capabilities of a single ant
- ant algorithms
  - study models derived from the observation of real ants behavior
  - use these models as a source of inspiration for the design of novel algorithm
  - for the solution of optimization and distributed control problems
- different aspects of the behavior $\rightarrow$different kinds of ant algorithms
  - aspects
    - foraging
    - division of labor
    - brood sorting
    - cooperative transport

Slide 5

From real to artificial ants (Cont'd)

- stigmergy
  - ants coordinate their activities via stigmergy
  - a form of indirect communication mediated by modifications of the environment.
  - e.g., a foraging ant deposits a chemical on the ground which increases the probability that other ants will follow the same path
  - from this, ants can achieve self-organization
- artificial stigmergy
  - the idea behind ant algorithms to coordinate societies of artificial agents
  - one of the most successful examples of ant algorithms $\rightarrow$ant colony optimization (ACO)
- ACO
  - inspired by the foraging behavior of ant colonies and targets discrete optimization problems

Ants' foraging behavior and optimization

- pheromones
  - chemicals produced by ants
  - a communication tool between individuals and the environment
  - trail pheromone
    - a specific type of pheromone
    - used for marking paths on the ground
    - foragers can follow the path to food discovered by other ants

Slide 6

- double bridge experiments
  - pheromone trail
    - while walking from food source to the nest and vice versa
    - ants deposit pheromones on the ground
    - forming in this way a pheromone trail
  - ants can smell the pheromone and they tend to choose, probabilistically, paths marked by strong pheromone concentrations
  - double bridge experiments

Ants' foraging behavior and optimization (Cont'd)

Slide 7

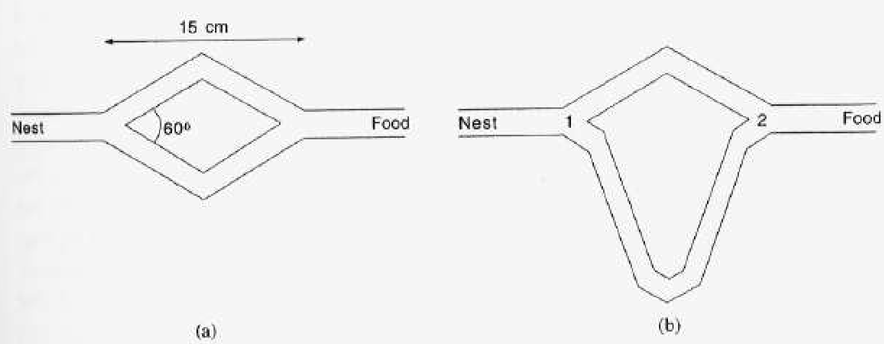

Figure 1.1
Experimental setup for the double bridge experiment. (a) Branches have equal length. (b) Branches have different length. Modified from Goss et al. (1989).

Ants' foraging behavior and optimization (Cont'd)

Slide 8

- first experiment
  - $r = l_l/l_s = 1$, where l_l is longer branch and l_s is shorter one
  - result: eventually all the ants used the same branch
    - when a trial starts, there is no pheromone
    - ants do not have a preference and select with the same probability
    - because random fluctuations, a few more ants will select one branch over the other
  - → more pheromone on that branch
  - → more ants to choose that branch
  - → finally, converge to one branch
  - this is called autocatalytic or positive feedback
  - macroscopic behavior & microscopic interactions
    - macroscopic behavior: convergence to one branch
    - microscopic interactions: local interactions among individuals via pheromone

Ants' foraging behavior and optimization (Cont'd)

Slide 9

- second experiment
  - $r = l_l/l_s = 2$
  - in most of the trials, all the ants chose to use only the short branch
  - initially, ants chose randomly (half choose the short and the other half the long)
  - the ants choosing the short branch are the first to reach the food and to start their return to the nest
  - ☛ pheromone starts to accumulate faster on the short branch
  - ☛ initial random fluctuations is much reduced, stigmergy, autocatalysis, and differential path length are the main mechanisms at work
  - not all the ants use the short branch, but a small percentage may take the longer one →path exploration

Ants' foraging behavior and optimization (Cont'd)

Slide 10

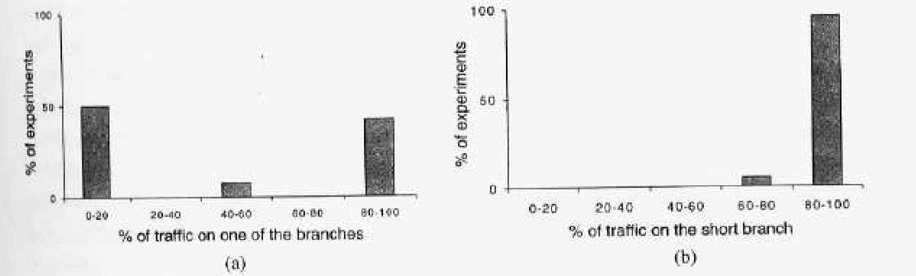

Figure 1.2
Results obtained with *Iridomyrmex humilis* ants in the double bridge experiment. (a) Results for the case in which the two branches have the same length ($r = 1$); in this case the ants use one branch or the other in approximately the same number of trials. (b) Results for the case in which one branch is twice as long as the other ($r = 2$); here in all the trials the great majority of ants chose the short branch. Modified from Goss et al. (1989).

Ants' foraging behavior and optimization (Cont'd)

– third experiment

Slide 11

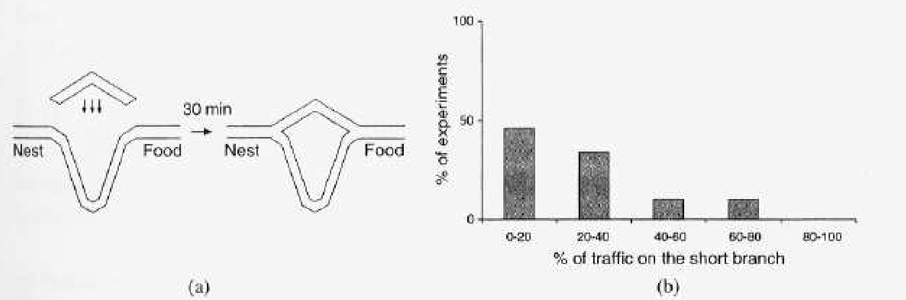

Figure 1.3
In this experiment initially only the long branch was offered to the colony. After 30 minutes, when a stable pheromone trail has formed on the only available branch, a new shorter branch is added. (a) The initial experimental setup and the new situation after 30 minutes, when the short branch was added. (b) In the great majority of the experiments, once the short branch is added the ants continue to use the long branch.

Ants' foraging behavior and optimization (Cont'd)

Slide 12

  * only long branch exists, after convergence (about 30 minutes), a new shorter branch adds
  * the short branch was only selected sporadically
  * the colony was trapped on the long branch
- ☛ by the high pheromone concentration on the long branch and by the slow evaporation of pheromone
- ☛ the pheromone evaporates too slowly to allow the ant colony to forget the suboptimal path

Ants' foraging behavior and optimization (Cont'd)

Slide 13

- a stochastic model
  - parameters
    * ψ ants per second
    * v cm/s
    * depositing one unit of pheromone
    * l_s, l_l short and long branch lengths
    * $t_s = l_s/v$ traverse time on short branch
    * $r = l_l/l_s \rightarrow r \cdot t_s$ traverse time on long branch
    * $p_{ia}(t)$ probability of selecting branch $a \in \{s, l\}$ at point $i \in \{1, 2\}$ (see fig. 1.1b)
    * $\varphi_{ia}(t)$ total amount of pheromone on the branch

Ants' foraging behavior and optimization (Cont'd)

Slide 14

– $p_{is}(t)$ probability of selecting short branch

$$p_{is}(t) = \frac{(t_s + \varphi_{is}(t))^\alpha}{(t_s + \varphi_{is}(t))^\alpha + (t_s + \varphi_{il}(t))^\alpha}$$

where the following constants

· $\alpha = 2$ was derived from experiments

· $p_{is}(t) + p_{il}(t) = 1$

· the amount of pheromone $\propto$ the number of ants

· no evaporation is considered

– differential equations

$$\begin{aligned} \mathrm{d}\varphi_{is}/\mathrm{dt} &= \psi p_{js}(t - t_s) + \psi p_{is}(t), (i{=}1,j{=}2;i{=}2,j{=}1), \\ \mathrm{d}\varphi_{il}/\mathrm{dt} &= \psi p_{jl}(t - r \cdot t_s) + \psi p_{il}(t), (i{=}1,j{=}2;i{=}2,j{=}1) \end{aligned}$$

Ants' foraging behavior and optimization (Cont'd)

Slide 15

– Monte Carlo simulation results of 1000 runs at $r = 1$ and $r = 2$

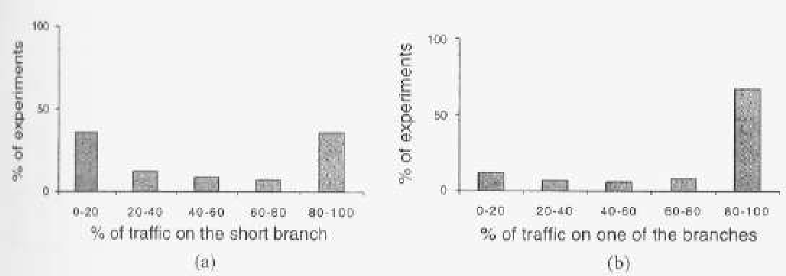

Figure 1.4
Results of 1000 Monte Carlo simulations of the model given by equations (1.1), (1.2), and (1.3), with $\psi = 0.5$ ant per second. Ants were counted between the 501st and 1000th ant crossing the bridge. (a) The ratio between the long and the short branch was set to $r = 1$. (b) The ratio between the long and the short branch was set to $r = 2$. Modified from Goss et al. (1989).

* when $r = 1$, ants select two branches with equal probability
* when $r = 2$, most of ants choose the short branch

Ants' foraging behavior and optimization (Cont'd)

Slide 16

- pheromone deposit
  - ants deposit pheromone both on their forward and backward paths
    if ants deposit pheromone only during the forward or only during backward trip
    then, the result is unstable to choose the shortest branch
  - this has confirmed by real ant experiments

Toward artificial ants

Slide 17

- graph model of double bridge

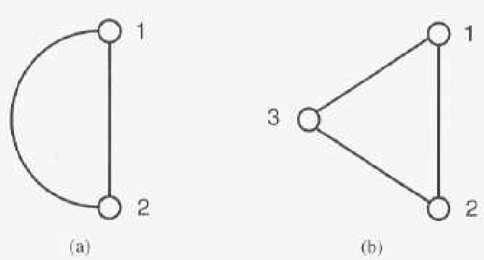

Figure 1.5
The graphs are two equivalent models of the experimental setup shown in figure 1.1b. In both cases, ants move from the nest to the food and back either via a short or via a long branch. (a) In this model the long branch is r times longer than the shorter one. An ant entering the long branch updates the pheromone on it r time units later. (b) In this model, each arc of the graph has the same length, and a longer branch is represented by a sequence of arcs. Here, for example, the long branch is twice as long as the short branch. Pheromone updates are done with one time unit delay on each arc.

Toward artificial ants (Cont'd)

Slide 18

- assumption
  - time is discrete
  - at each time step, each ant moves toward a neighbor node at constant speed of one unit of length
  - ants add one unit of pheromone to the arcs
- short and long path probability

$$p_{is}(t) = \frac{[\varphi_{is}(t)]^\alpha}{[\varphi_{is}(t)]^\alpha + [\varphi_{il}(t)]^\alpha},\ p_{il}(t) = \frac{[\varphi_{il}(t)]^\alpha}{[\varphi_{is}(t)]^\alpha + [\varphi_{il}(t)]^\alpha}$$

- trail update

$$\varphi_{is}(t) = \varphi_{is}(t-1) + p_{is}(t-1)m_i(t-1) + p_{js}(t-1)m_j(t-1), (i{=}1, j{=}2; i{=}2, j{=}1)$$

$$\varphi_{il}(t) = \varphi_{il}(t-1) + p_{il}(t-1)m_i(t-1) + p_{jl}(t-r)m_j(t-r), (i{=}1, j{=}2; i{=}2, j{=}1)$$

where $m_i(t) = p_{js}(t-1)m_j(t-1) + p_{jl}(t-r)m_j(t-r), (i{=}1, j{=}2; i{=}2, j{=}1)$

Toward artificial ants (Cont'd)

Slide 19

- difference of previous model
  - it considers the average behavior of the system, not the stochastic behavior of the single ant
  - it is a discrete time model
- simulation run is very similar to that of the continuous time model

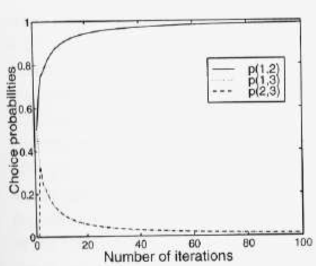

Figure 1.6
Result of the simulation of the model described by equations (1.4) through (1.7). The figure shows the probability of choosing the three branches of the graph in figure 1.5b. After a short transitory period the probabilities of choosing the long branch (($p(1,3) \equiv p_{1l}$ and ($p(2,3) \equiv p_{2l}$) become vanishingly small (in the graph they are superimposed after a few iterations from the start), while the probability of choosing the short branch ($p(1,2) \equiv p_{1s} \equiv p_{2s}$) tends to 1. Note that probabilities are symmetric: $p(i,j) = p(j,i)$. Parameter settings: $\alpha = 2$, $r = 2$, $t = 100$.

Artificial ants and minimum cost paths

Slide 20

- goal
  - define algorithms to solve minimum cost problems on more complicated graphs

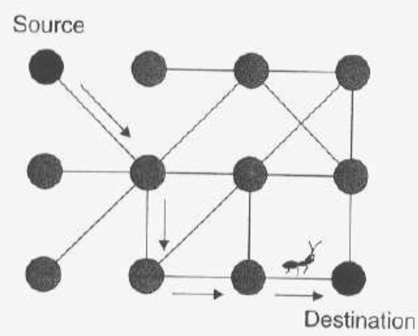

Figure 1.7
Ants build solutions, that is, paths from a source to a destination node.

Artificial ants and minimum cost paths (Cont'd)

Slide 21

- problems using the artificial ants mentioned previous
  - loop
    * loops tend to become more and more attractive by the forward pheromone trail updating mechanism
    * even if an ant can escape such loops, short paths are no longer favored
  - the simplest solution
    * removal of the forward updating mechanism

    → cause the algorithm to be unstable as mentioned before

Artificial ants and minimum cost paths (Cont'd)

Slide 22

- solution via the use of memory
  - ∗ probabilistic solution construction biased by pheromone trails without forward pheromone updating
  - ∗ deterministic backward path with loop elimination and with pheromone updating
  - ∗ evaluation of the quality of the solutions generated and use of the solution quality in determining the quantity of pheromone to deposit
  - √ pheromone evaporation greatly improved
  - ☛ Simple ACO (S-ACO)
- probabilistic forward ants and solution construction
  - ∗ forward mode: nest →food, backward mode: food →nest
  - ∗ in forward mode ants build a solution by choosing probabilistically the next node
  - ∗ the probabilistic choice is biased by pheromone trails previously deposited by other ants
  - ∗ forward ants do not deposit any pheromone

Artificial ants and minimum cost paths (Cont'd)

Slide 23

- deterministic backward ants and pheromone update
  - ∗ an ant can retrace the path it has followed
  - ∗ ants improve the system performance by implementing loop elimination
  - ∗ S-ACO ants leave pheromone on the arcs they traverse
- pheromone updates based on solution quality
  - ∗ at destination (food), an ant evaluates the cost of the solution
  - ∗ use this evaluation to modulate the amount of pheromone they deposit in backward mode
  - → more strongly toward better solutions
- pheromone evaporation
  - ∗ in real ant, pheromone intensity decreases over time because of evaporation
  - ∗ in S-ACO, pheromone evaporation rule is applied
  - ∗ very useful in S-ACO

Artificial ants and minimum cost paths (Cont'd)

Slide 24

- S-ACO
  - parameters
    - ∗ a graph $G = (N, A)$
    - ∗ τ_{ij}: artificial pheromone trail on arc (i, j)
    - ∗ initial pheromone: a constant amount of pheromone (e.g. $\tau_{ij} = 1, \forall (i, j) \in A$)
  - ants' path-searching behavior
    - ∗ probability that an ant k at a node i choose j node

$$p_{ij}^k = \begin{cases} \frac{\tau_{ij}^\alpha}{\sum_{l \in N_i^k} \tau_{il}^\alpha}, & \text{if } j \in N_i^k; \\ 0, & \text{if } j \notin N_i^k; \end{cases}$$

where N_i^k is the neighborhood of ant k in node i.

- ∗ N_i^k is all nodes directly connected to node i except for the predecessor of node i (last visited node)
- ∗ only in case N_i^k is empty, node i's predecessor is included into N_i^k
- ∗ ants hop nodes until it reaches the destination node

→ can easily lead to loops

Artificial ants and minimum cost paths (Cont'd)

Slide 25

- path retracing and pheromone update
  - ∗ when reaching the dest. node, the ant switches from the forward mode to the backward mode
  - ∗ then retraces step by step the same path backward to the source node
  - ∗ before starting the return trip, an ant eliminates the loops
  - ∗ loop elimination
    - · scan the nodes from the source
    - · for ith node, the path is scanned starting from the dest. node until the first occurrence of the node is encountered at j node
    - · the subpath from $i + 1$ to j, a loop, can be eliminated

Artificial ants and minimum cost paths (Cont'd)

Slide 26

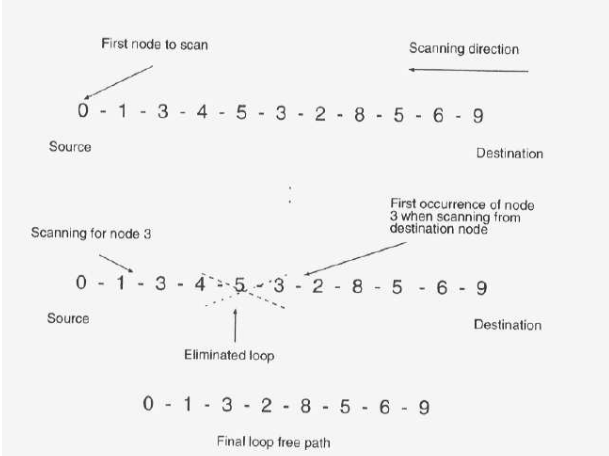

Figure 1.8
Illustration of the scanning process for loop elimination.

· for example, the loop 3-4-5-3 of length 3 is eliminated, but not guaranteed the longest 5-3-2-8-5 loop elimination

Artificial ants and minimum cost paths (Cont'd)

Slide 27

∗ during retrace, the ant deposits an amount $\Delta\tau^k$ of pheromone on arc (i, j):
$\tau_{ij} \leftarrow \tau_{ij} + \Delta\tau^k$

∗ decision of $\Delta\tau_k$

· in the simplest case, constant value can be used

→ only the differential path length works (ants on a shorter path deposit earlier than ants on a longer path)

· a function of the path length can be used (the shorter the path the more pheromone is deposited)

Artificial ants and minimum cost paths (Cont'd)

Slide 28

- pheromone trail evaporation
  - ∗ an exploration mechanism that avoids quick convergence of all the ants toward a suboptimal path
  - ∗ in artificial ants, it is important
  - ∗ in real ants, pheromone trails also evaporate, but not play an important role why? →it seems because the path that real ants find is very simple
  - ∗ evaporation

$$\tau_{ij} \leftarrow (1-\rho)\tau_{ij}, \forall(i,j) \in A$$

    where $\rho \in (0,1]$ is a parameter
- an iteration of S-ACO
  - ∗ ant's movement
  - ∗ pheromone evaporation
  - ∗ pheromone deposit

Artificial ants and minimum cost paths (Cont'd)

Slide 29

- experiments with S-ACO
  - two simple graphs
    - ∗ double bridge (fig. 1.5b)
    - ∗ extended double bridge (fig. 1.9)

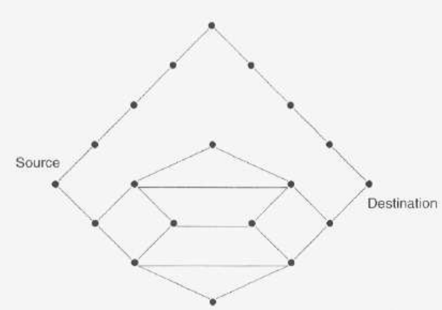

Figure 1.9
Extended double bridge. An ant starting in the source node can choose between the upper and the lower parts of the graph. The upper part consists of a single path of length 8 leading directly to the destination node. Differently, the lower part of the graph consists of a set of paths (which includes many paths shorter than eight steps) and the ant has many decisions to do before reaching the destination node. Therefore, ants choosing the upper part will always find a path of length 8, while ants choosing the lower part may find paths shorter than 8, but they may also enter loops and generate very long paths.

Artificial ants and minimum cost paths (Cont'd)

Slide 30

· one loop-free, but worse than optimal, upper path of the graph
· lower part of the graph that contains two optimal paths of length 5
· many longer loop-free paths and infinitely many, much longer "loopy" paths

– number of ants and types of pheromone update
* experiments with double bridge (fig. 1.5b)
* comparing the results of S-ACO to those of real ants' behavior (artificial ants)
* two experiments
· run S-ACO with different number of ants and with depositing a constant of pheromone
· same as above except that ants deposit an amount of pheromone that is a inversely proportional to the length of the path (i.e., $\Delta\tau_k = 1/L^k$, where L^k is the length of ant k's path)
* parameters
· 100 trials
· stop after each ant had moved 1000 steps
· evaporation was set to $\rho = 0$
· the parameter $\alpha = 2$

Artificial ants and minimum cost paths (Cont'd)

Slide 31

* results

Table 1.1
Percentage of trials in which S-ACO converged to the long path (100 independent trials for varying values of m, with $\alpha = 2$ and $\rho = 0$)

m	1	2	4	8	16	32	64	128	256	512
without path length	50	42	26	29	24	18	3	2	1	0
with path length	18	14	8	0	0	0	0	0	0	0

Column headings give the number m of ants in the colony. The first row shows results obtained performing pheromone updates without considering path length; the second row reports results obtained performing pheromone updates proportional to path length.

· percentage that pheromone trail was higher on the long path
· without path length
up to 32 ants, S-ACO converged relatively often to the longer path
→ due to probabilistic selection of initial path and strong reinforcement
→ drastically decreased above 64
· with path length
solution quality are much better than those without path length
8 ants or more always converged to the short path

Artificial ants and minimum cost paths (Cont'd)

Slide 32

- influence of the parameter α
  * changed in step sizes of 0.25 from 1 to 2
  * results
    · without path length, increasing α had a negative effect on the convergence because large values of α tend to amplify the influence of initial random fluctuations
    · with path length, the results were rather independent of the particular value of α
  ☛ effectiveness of the differential path length effect strongly decreases with increasing problem complexity
- pheromone evaporation
  * experiments with extended double bridge
  * do loop elimination
  * pheromone deposit: $\Delta\tau^k = 1/L^k$
  * plot the moving averages using the $4 \cdot m$ most recent paths found by the ants
  * experiments with the evaporation rate $\rho \in \{0, 0.01, 0.1\}$, $\alpha = 1$, and $m = 128$

Artificial ants and minimum cost paths (Cont'd)

Slide 33

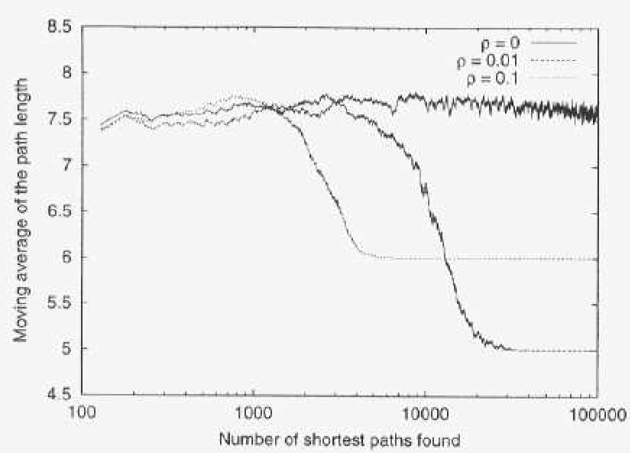

Figure 1.10
The graph plots the moving averages (given on the y-axis) of the ants' path length for the graph of figure 1.9 as a function of the number of completed paths (given on the x-axis). We give plots for not using evaporation ($\rho = 0$), low evaporation ($\rho = 0.01$), and high evaporation ($\rho = 0.1$). The trials were stopped after 5000 iterations; $\alpha = 1$ and $m = 128$.

Artificial ants and minimum cost paths (Cont'd)

Slide 34

- results
  - $\cdot$ no evaporation rate $\rho = 0$, the algorithm does not converge
  - $\cdot$ with pheromone evaporation, the algorithm converges either to the value 5 for $\rho = 0.01$ or to the value 6 for $\rho = 0.1$
- general observations
  - $\cdot$ without pheromone updates based on solution quality, S-ACO is much worse,
  - $\rightarrow$ the larger the parameters α or ρ, the faster S-ACO converges to the suboptimal solution of length 8
  - $\cdot$ pheromone evaporation rate ρ can be critical
  - $\rightarrow$ if ρ is set to high, S-ACO often converged to suboptimal paths
  - $\rightarrow$ for examples, when $\rho = 0.2$ in 15 trials, once to a path of length 8, once to a path of length 7, twice to a path of length 6 but for $\rho = 0.01$, S-ACO converged to the shortest path in all trials
  - $\cdot$ large value of α generally result in a worse behavior because of initial random fluctuations

Artificial ants and minimum cost paths (Cont'd)

Slide 35

- discussion
  - for more complex problems, the parameter setting of S-ACO become increasingly important to converge the optimal solution
  - conclusions
    - $\cdot$ the differential path length effect is not enough for large problems
    - $\cdot$ pheromone updates based on solution quality are important for fast convergence
    - $\cdot$ large α lead to a strong emphasis of initial, random fluctuations
    - $\cdot$ the larger the number of ants, the better the convergence behavior, but need longer simulation times
    - $\cdot$ pheromone evaporation is important when trying to solve more complex problems

Ant colony optimization algorithms for the TSP

Slide 36

- traveling salesman problem (TSP)
  - good for testing because
    - an extensively studied problem
    - an important NP-hard optimization problem
    - a standard test bed for new algorithmic ideas
    - ACO algorithms are easily applied
  - informal description
    - start from a salesman hometown
    - find a shortest tour of customer cities, visiting each customer city exactly once
    - then back home
  - formal description
    - a complete weighted graph $G = (N, A)$ with N cities and A arcs
    - each arc $(i, j) \in A$ is assigned a value (length) d_{ij} between cities i and j with $i, j \in N$

Ant colony optimization algorithms for the TSP (Cont'd)

Slide 37

  - types
    - asymmetric TSP: there is at least one arc (i, j) for which $d_{ij} \neq d_{ji}$
    - symmetric TSP: $d_{ij} = d_{ji}$ for all arcs $(i, j) \in A$
  - goal
    - find a minimum length Hamiltonian circuit of the graph
    - a Hamiltonian circuit is a closed path visiting each of $n = |N|$ nodes of G exactly once
  - an optimal solution
    - a permutation π of the node indices $\{1, 2, \ldots, n\}$ such that the length $f(\pi)$ is minimal, where

$$f(\pi) = \sum_{i=1}^{n-1} d_{\pi(i)\pi(i+1)} + d_{\pi(n)\pi(1)}$$

Slide 38

Ant colony optimization algorithms for the TSP (Cont'd)

- ACO algorithms for the TSP
  - parameters
    - τ_{ij}: desirability of visiting city j directly after city i
    - $\eta_{ij} = 1/d_{ij}$: heuristic desirability of going from city i directly to city j

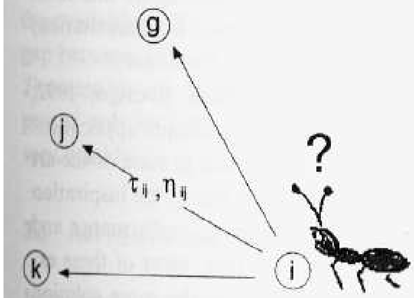

Figure 3.2
An ant arriving in city i chooses the next city to move to as a function of the pheromone values τ_{ij} and of the heuristic values η_{ij} on the arcs connecting city i to the cities j the ant has not visited yet.

Slide 39

Ant colony optimization algorithms for the TSP (Cont'd)

- high level description of ACO
  - tour construction
    - choose a start city
    - use pheromone and heuristic values to probabilistically construct a tour by iteratively adding cities that the ant has not visited yet, until all cities have been visited
    - go back to the initial city
  - pheromone deposit
    - after all ants have completed their tour, they may deposit pheromone on the tours
  - local search
    - before adding pheromone, the tours may be improved by the application of local search

Ant colony optimization algorithms for the TSP (Cont'd)

Slide 40

– ACO algorithms

```
procedure ACOMetaheuristicStatic
  Set parameters, initialize pheromone trails
  while (termination condition not met) do
    ConstructAntsSolutions
    ApplyLocalSearch
    UpdatePheromones
  end
end
```

– extensions
  * first ACO algorithm: ant system (AS)
  * based on AS
    · elitist AS
    · rank-based AS
    · MAX-MIN AS

Ant colony optimization algorithms for the TSP (Cont'd)

Slide 41

  * main differences between AS and extensions are
    · the way the pheromone update is performed
    · the management of the pheromone trails
  * substantial modification
    · Ant-Q
    · Ant Colony System (ACS)
    · ANTS algorithm
    · Hyper-cube framework for ACO

Ant system and its direct successors

Slide 42

- Ant System
  - three different versions of AS were proposed
    - ant-density and ant-quantity
      - update pheromone directly after a move from one city to an adjacent city
      - were abandoned because of their inferior performance
    - ant-cycle
      - update pheromone after all the ants had constructed the tours
      - the amount of pheromone deposited $\propto$ the tour quality
  - two main phases
    - the ants' solution construction
    - the pheromone update

Ant system and its direct successors (Cont'd)

Slide 43

- a good heuristic to initialize the pheromone trail
  - set them to a value slightly higher than the expected amount of pheromone deposited by the ants in one iteration
  - a rough estimate of this value, $\forall(i,j), \tau_{ij} = \tau_0 = m/C^{nn}$, where m is the number of ants and C^{nn} is the length of a tour generated by the nearest-neighbor heuristic
  - if τ_0 is too low, then the search is quickly biased by the first tours
  - if τ_0 is too high, many iterations are lost until pheromone evaporation

Ant system and its direct successors (Cont'd)

Slide 44

- tour construction
  - initially, ants are put on randomly chosen cities
  - ant k applies a probabilistic action choice rule to decide next city

$$p_{ij}^k = \frac{[\tau_{ij}]^\alpha [\eta_{ij}]^\beta}{\sum_{l \in N_i^k} [\tau_{il}]^\alpha [\eta_{il}]^\beta}, \quad \text{if } j \in N_i^k$$

    where
    - · $\eta_{ij} = 1/d_{ij}$ is a heuristic value
    - · α and β are two parameters of relative influence
    - · N_i^k is the feasible neighborhood of ant k
  - if $\alpha = 0$
    - · the closest cities are more likely to be selected
    - · corresponds to a classic stochastic greedy algorithm
  - if $\beta = 0$
    - · only pheromone amplification is at work
    - · leads to rather poor results
    - · for $\alpha > 1$, leads to a stagnation situation (all ants follow the same path)

Ant system and its direct successors (Cont'd)

Slide 45

- a memory M^k
  - · contain the cities already visited in the order
  - · used to define the feasible neighborhood N_i^k
  - · used to compute the length of the tour T^k and to retrace the path to deposit pheromone
- two different ways of solution construction
  - · sequential: an ant builds a complete tour before the next one starts
  - · parallel: all the ants move from their current city to the next one

Ant system and its direct successors (Cont'd)

Slide 46

- update of pheromone trails
  - pheromone evaporation

$$\tau_{ij} \leftarrow (1-\rho)\tau_{ij}, \quad \forall (i,j) \in L$$

    where $0 < \rho \leq 1$ is the pheromone evaporation rate
  - ρ
    - used to avoid unlimited accumulation of pheromone
    - used to forget bad decisions previously taken
  - pheromone trail

$$\tau_{ij} \leftarrow \tau_{ij} + \sum_{k=1}^{m} \Delta\tau_{ij}^{k}, \quad \forall (i,j) \in L$$

    where $\Delta\tau_{ij}^{k}$ is the amount of pheromone ant k deposits on the arcs it has visited
  - $\Delta\tau_{ij}^{k}$

$$\Delta\tau_{ij}^{k} = \begin{cases} 1/C^{k}, & \text{if arc}(i,j) \text{ belongs to } T^{k}; \\ 0, & \text{otherwise;} \end{cases}$$

    where C^{k} is the length of the tour T^{k}

Ant system and its direct successors (Cont'd)

Slide 47

- AS behavior
  - the better an ant's tour is, the more pheromone the arcs belongs to this tour receive
  - arcs used by many ants and which are part of short tours, receive more pheromone
  - those arcs are more likely to be chosen in future iterations

Ant system and its direct successors (Cont'd)

Slide 48

- elitist ant system (EAS)
  - idea
    - provide strong additional reinforcement to the best tour T^{bs} (*best-so-far* tour)
  - update of pheromone trails
    - add a quantity e/C^{bs} to the tour T^{bs}
    - e is a parameter that defines the weight of T^{bs} and C^{bs} is its length

$$\tau_{ij} \leftarrow \tau_{ij} + \sum_{k=1}^{m} \Delta\tau_{ij}^{k} + e\Delta\tau_{ij}^{bs},$$

where

$$\Delta\tau_{ij}^{bs} = \begin{cases} 1/C^{bs}, & \text{if arc}(i,j) \text{ belongs to } T^{bs}; \\ 0, & \text{otherwise;} \end{cases}$$

  - results
    - find better tours and find them in a lower number of iterations with an appropriate value of e

Ant system and its direct successors (Cont'd)

Slide 49

- rank-based ant system (AS$_{rank}$)
  - update of pheromone trails
    - deposit an amount of pheromone that decreases with its rank
    - only the $(w-1)$ best ranked ants and the best-so-far tour ant (not necessarily in current iteration) deposit pheromone

$$\tau_{ij} \leftarrow \tau_{ij} + \sum_{r=1}^{w-1} (w-r)\Delta\tau_{ij}^{r} + w\Delta\tau_{ij}^{bs},$$

where $\Delta\tau_{ij}^{r} = 1/C^{r}$ and $\Delta\tau_{ij}^{bs} = 1/C^{bs}$

  - results
    - slightly better than EAS and significantly better than AS

Ant system and its direct successors (Cont'd)

Slide 50

- MAX-MIN AS
  - four main modifications
    * only either the iteration-best or best-so-far ant can deposit pheromone →lead to a stagnation situation
    * limit the possible range of pheromone trail values to the interval $[\tau_{min}, \tau_{max}]$
    * initialize to the upper pheromone trail limit with a small evaporation rate →increase the exploration of tours at the start
    * reinitialize each time the system approaches stagnation or when no improved tour has been generated for a certain number of consecutive iterations
  - update of pheromone trails

$$\tau_{ij} \leftarrow \tau_{ij} + \Delta\tau_{ij}^{best},$$

where $\Delta\tau_{ij}^{best} = 1/C^{best}$, C^{best} can be either the best-so-far C^{bs} or the iteration-best C^{ib}

    * the choice of the relative frequency with two rules influence on how greedy the search is

Ant system and its direct successors (Cont'd)

Slide 51

  - results
    * for small TSPs, only iteration-best pheromone update is best
    * for large TSPs (several hundreds of cities), emphasis to the best-so-far update is better
    * need gradually increasing the frequency of the best-so-far update as the cities increase
  - pheromone trail limits
    * $\tau_{max} = 1/\rho C^{bs}$, $\tau_{min} = \tau_{max}/a$ where a is a parameter
    * in order to avoid stagnation, τ_{min} plays a more important role than τ_{max}
    * τ_{max} is useful for trail reinitializations

Extensions of ant system

Slide 52

- two additional ACO algorithms
  - strongly inspired by AS
  - introduction of new mechanisms based on ideas not included in original AS
  - extensions
    * ant colony system
    * approximate nondeterministic tree search (ANTS)
- ant colony system (ACS)
  - three main points
    * exploit the search experience more strongly than AS
    * pheromone evaporation and pheromone deposit take place only on the best-so-far tour
    * each time an ant uses an arc (i, j) to move from city i to city j, it removes some pheromone to increase the exploration of alternative paths

Extensions of ant system (Cont'd)

Slide 53

- tour construction
  * pseudorandom proportional rule

$$j = \begin{cases} \text{argmax}_{l \in N_i^k} \{\tau_{il}[\eta_{il}]^\beta\}, & \text{if } q \leq q_0; \\ J, & \text{otherwise}; \end{cases}$$

where
· q is a random variable uniformly distributed in $[0, 1]$
· q_0 is a parameter
· J is a random variable selected according to the probability of AS with $\alpha = 1$

$$p_{ij}^k = \frac{[\tau_{ij}]^\alpha [\eta_{ij}]^\beta}{\sum_{l \in N_i^k} [\tau_{il}]^\alpha [\eta_{il}]^\beta}, \quad \text{if } j \in N_i^k$$

- operation
  · with probability q_0, the ant selects the best possible move by the learned pheromone trails and the heuristic info.
  · with probability $(1 - q_0)$, the ant performs a biased exploration

Extensions of ant system (Cont'd)

Slide 54

- global pheromone trail update
  - only the best-so-far ant is allowed to add pheromone after each iteration

$$\tau_{ij} \leftarrow (1-\rho)\tau_{ij} + \rho\Delta\tau_{ij}^{bs}, \quad \forall(i,j) \in T^{bs}$$

    where $\Delta\tau_{ij}^{bs} = 1/C^{bs}$
  - operation
    - pheromone evaporation and deposit apply only to the arcs of T^{bs}, not to all
    - computational complexity is reduced from $O(n^2)$ to $O(n)$
- experimental results
  - for small TSP, the best-so-far and the iteration-best tour shows similar performances
  - for large TSP (more than 100 cities), the use of the best-so-far gives far better results

Extensions of ant system (Cont'd)

Slide 55

- local pheromone trail update
  - apply immediately after having crossed an arc (i,j) during the tour construction

$$\tau_{ij} \leftarrow (1-\xi)\tau_{ij} + \xi\tau_0,$$

    where
    - $\xi, 0 < \xi < 1$, and τ_0 are two parameters
    - τ_0 is set to be the same as the initial value for the pheromone trails
    - experimentally, a good value for ξ is 0.1 and $\tau_0 = 1/nC^{nn}$, where n is the number of cities and C^{nn} is the length of a nearest-neighbor tour
  - effect
    - the local updating rule reduces pheromone trail τ_{ij}
      $\rightarrow$less desirable for the following ants
      $\rightarrow$increase in the exploration of arcs that have not been visited yet
      $\rightarrow$no stagnation behavior

Extensions of ant system (Cont'd)

Slide 56

- some additional remarks
  * first ACO algorithm to use *candidate lists*
  * to restrict the number of available choices at each construction step
  * candidate lists
    · a number of the best rated choices according to some heuristic criterion
    · an example
    → sort the neighbors of a city i according to nondecreasing distances
    → insert a fixed number *cand* of closest cities into i's candidate list (remaining fixed during process)
    → ants select the next city among i's candidate list not visited yet (only if all the cities of the candidate list are already marked as visited, one of the remaining cities is chosen)
    · results
    → the use of candidate lists improves the solution quality and leads to a significant speedup

behavior of ACO algorithms

Slide 57

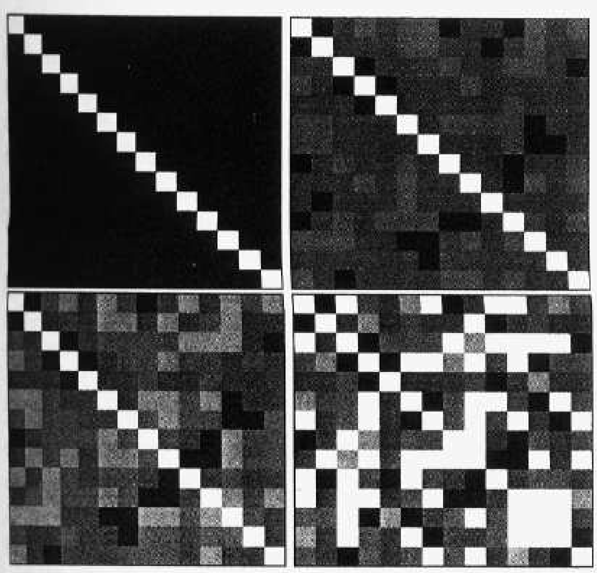

Figure 3.5
A visual representation of the pheromone matrix. The pheromone values on the arcs, stored in the pheromone matrix, are translated into gray-scale values; the darker an entry, the higher the associated pheromone trail value. The plots, from upper left to lower right, show the pheromone value for AS applied to TSPLIB instance `burma14` with 14 cities after 0, 5, 10, and 100 iterations. Note the symmetry with respect to the main diagonal, which is due to the fact that `burma14` is a symmetric TSP instance.

behavior of ACO algorithms (Cont'd)

Slide 58

- behavior
  - at the beginning, all the matrix's cells are black (high pheromone) except for diagonal's cells (zero pheromone)
  - after five iterations, the differences between pheromone trails are low
  - after five iterations, the differences are rather high
  - at 100 iterations, the differences become stronger and finally a situation that only few connections have a large amount of pheromone reaches
- stagnation
  - with good parameters, iterations progressively reduce the size of the explored search space
  - this behavior may become undesirable, early stagnation of the search
  - → cease to explore new possibilities and no better tour is likely to be found anymore

behavior of ACO algorithms (Cont'd)

Slide 59

- several measures to detect
  - ∗ compute the standard deviation σ_L of the length of the tours
    - · if σ_L is zero, all ants follow the same path (unlikely different path of the same length)
    - · this measure depends on the absolute values of the tour lengths
      →a better choice: $\sigma_L/\bar{L}$
  - ∗ compute the distance $dist(T, T^{'})$ between two tours T and $T^{'}$ (count the number of arcs contained only one tour)
    - · if average distance decreases, preferred paths are appearing
    - · if the average distance becomes zero, then search stagnation
      →computationally expensive $O(n^2)$
  - ∗ λ-branching factor, $0 < \lambda < 1$
    - · the number of arcs having $\tau_{ij} \geq \tau^i_{min} + \lambda(\tau^i_{max} - \tau^i_{min})$
    - · range over the interval $[2, n-1]$, where n is the number of cities
    - · if $\bar{\lambda}$ is very close to 3, only three arcs have a high probability of being chosen (in TSP, the minimal $\bar{\lambda}$ is 2)
      →difficult to choose the λ

behavior of ACO algorithms (Cont'd)

Slide 60

∗ average entropy

$$\bar{\varepsilon} = \sum_{i=1}^{n} \varepsilon_i / n, \quad \varepsilon_i = -\sum_{j=1}^{l} p_{ij} \log p_{ij},$$

where p_{ij} is the probability of choosing arc (i, j), l is the number of possible choices, $1 \leq l \leq n - 1$

∗ a measure

$$\frac{\sum_{\tau_{ij} \in T} \min\{\tau_{max} - \tau_{ij}, \tau_{ij} - \tau_{min}\}}{n^2}$$

whose value tends to 0 as the algorithm moves toward stagnation

behavior of ACO algorithms (Cont'd)

Slide 61

- behavior of AS
  - behavior of the average λ-branching factor and of the average distance
  - two parameters (good : $\alpha = 1,\ \beta = 2,\ m = n$)(bad : $\alpha = 5,\ \beta = 0,\ m = n$)
  - bad behavior by early stagnation ($\lambda = 0.05$)

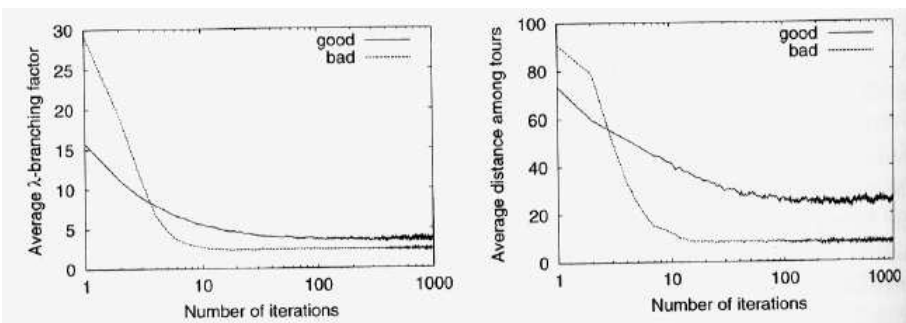

Figure 3.6
Bad behavior because of early stagnation: The plots give (left) the average λ-branching factor, $\lambda = 0.05$, and (right) the average distance among the tours generated by AS on the symmetric TSPLIB instance kroA100. "Good" system behavior is observed setting parameters to $\alpha = 1$, $\beta = 2$, $m = n$. "Bad" system behavior is observed setting parameters $\alpha = 5$, $\beta = 0$, $m = n$.

behavior of ACO algorithms (Cont'd)

Slide 62

- two parameters (good : $\alpha = 1,\ \beta = 2,\ m = n$)(bad : $\alpha = 1,\ \beta = 0,\ m = n$)
- bad behavior by excessive exploration ($\lambda = 0.05$)

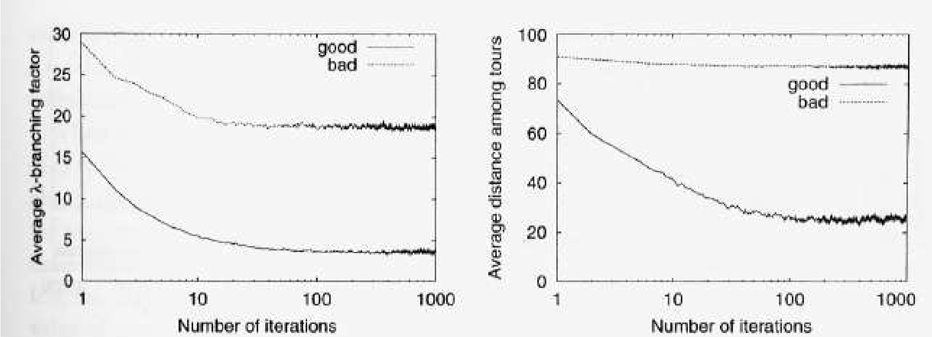

Figure 3.7
Bad behavior because of excessive exploration: The plots give (left) the average λ-branching factor, $\lambda = 0.05$, and (right) the average distance among the tours generated by AS on the symmetric TSPLIB instance kroA100. "Good" system behavior is observed setting parameters to $\alpha = 1, \beta = 2, m = n$. "Bad" system behavior is observed setting parameters $\alpha = 1,\ \beta = 0,\ m = n$.

behavior of ACO algorithms (Cont'd)

Slide 63

- results
  - ∗ a reasonable balance needs between a too narrow focus (cause stagnation) and a too weak guidance (cause excessive exploration)

behavior of ACO algorithms (Cont'd)

Slide 64

- behavior of extensions of AS
  - behavior of the average λ-branching factor and of the average distance

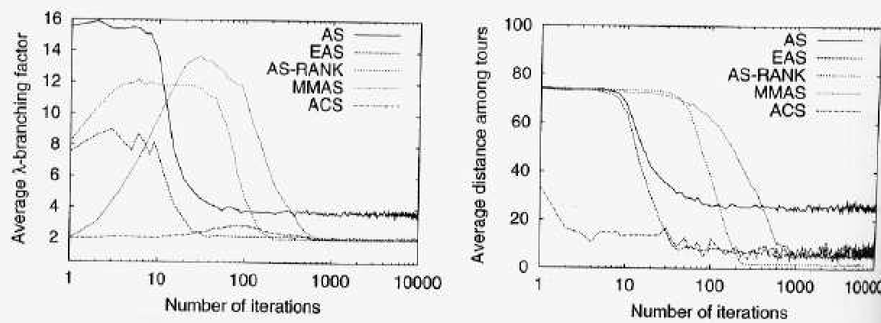

Figure 3.8
Comparing AS extensions: The plots give (left) the average λ-branching factor, $\lambda = 0.05$, and (right) the average distance among the tours for several ACO algorithms on the symmetric TSPLIB instance kroA100. Parameters were set as in box 3.1, except for β which was set to $\beta = 2$ for all the algorithms.

behavior of ACO algorithms (Cont'd)

Slide 65

- results
  - ∗ ACS shows a low λ-branching factor and small average distances
    - · explorative phase: high average λ-branching
    - · exploitation phase: very low average λ-branching
  - ∗ AS and AS$_{rank}$ show very soon transition from explorative phase to exploitation phase
  - ∗ MAX-MIN AS and AS$_{rank}$ occur later
    - · MAX-MIN AS needs occasional pheromone trail reinitializations
    - · AS$_{rank}$ could profit from occasional pheromone trail reinitializations

behavior of ACO algorithms (Cont'd)

Slide 66

- MAX-MIN AS and ACS
  - the most used and often the best-performing of ACO algorithms
- behavior of MAX-MIN AS
  - show exploration when it converges
    - converge to the minimum average λ-branching, stagnation
    - average distance remains high because of using τ_{max} and τ_{min}
    - τ_{min} provides to each arc a minimum probability $p_{min} > 0$
  - show the longest explorative search phase due to the fact
    - pheromone trails are initialized to the estimate of τ_{max}
    - evaporation rate is set to a low value ($\rho = 0.02$)
- behavior of ACS
  - focus around the best-so-far tour T^{bs} (for large value of q_0)
  - exploration is still possible because the pheromone of traversed arcs is diminished, making them less attractive