

Chapter 9

THE ANT COLONY OPTIMIZATION METAHEURISTIC: ALGORITHMS, APPLICATIONS, AND ADVANCES

Marco Dorigo

IRIDIA, Université Libre de Bruxelles, Belgium

E-mail: mdorigo@ulb.ac.be

URL: <http://iridia.ulb.ac.be/~mdorigo>

Thomas Stützle

Intellectics Group, TU Darmstadt, Germany

E-mail: stuetzle@informatik.tu-darmstadt.de

URL: <http://www.intellektik.informatik.tu-darmstadt.de/~tom>

1 INTRODUCTION

Ant Colony Optimization (ACO) [32,33] is a recent metaheuristic approach for solving hard combinatorial optimization problems. The inspiring source of ACO is the pheromone trail laying and following behavior of real ants which use pheromones as a communication medium. In analogy to the biological example, ACO is based on the indirect communication of a colony of simple agents, called (artificial) ants, mediated by (artificial) pheromone trails. The pheromone trails in ACO serve as a distributed, numerical information which the ants use to probabilistically construct solutions to the problem being solved and which the ants adapt during the algorithm's execution to reflect their search experience.

The first example of such an algorithm is Ant System (AS) [30,36–38], which was proposed using as example application, the well known Traveling Salesman Problem (TSP) [60,76]. Despite encouraging initial results, AS could not compete with state-of-the-art algorithms for the TSP. Nevertheless, it had the important role of stimulating further research both on algorithmic variants, which obtain much better computational performance, and on applications to a large variety of different problems. In fact, there exists now a considerable number of applications obtaining world class performance on problems including the quadratic assignment, vehicle routing, sequential ordering, scheduling, routing in Internet-like networks, and so on [22,26,46,47,67,85]. Motivated by this success, the ACO metaheuristic has been proposed [32,33] as a common framework for the existing applications and algorithmic variants. Algorithms which follow the ACO metaheuristic will be called in the following ACO algorithms.

Current applications of ACO algorithms fall into the two important problem classes of static and dynamic combinatorial optimization problems. Static problems are those whose topology and costs do not change while the problems are being solved. This is the case, e.g., for the classic TSP, in which city locations and intercity distances do not change during the algorithm's run-time. In contrast, in dynamic problems the topology and costs can change while solutions are built. An example of such a problem is routing in telecommunications networks [26], in which traffic patterns change all the time. The ACO algorithms for solving these two classes of problems are very similar from a high-level perspective, but they differ significantly in implementation details. The ACO metaheuristic captures these differences and is general enough to comprise the ideas common to both application types.

The (artificial) ants in ACO implement a randomized construction heuristic which makes probabilistic decisions as a function of artificial pheromone trails and possibly available heuristic information based on the input data of the problem to be solved. As such, ACO can be interpreted as an extension of traditional construction heuristics which are readily available for many combinatorial optimization problems. Yet, an important difference with construction heuristics is the adaptation of the pheromone trails during algorithm execution to take into account the cumulated search experience.

The rest of this chapter is organized as follows. In Section 2, we briefly overview construction heuristics and local search algorithms. In Section 3, we define the type of problem to which the ACO metaheuristic applies, the ants' behavior, and the ACO metaheuristic itself. Section 4 outlines the inspiring biological analogy and describes the historical developments leading to ACO. In Section 5, we illustrate how the ACO metaheuristic can be applied to different types of problems and we give an overview of its successful applications. Section 6 discusses several issues arising in the application of the ACO metaheuristic; Section 7 reports on recent developments and in Section 8 we conclude indicating future research directions.

2 TRADITIONAL APPROXIMATION APPROACHES

Many important combinatorial optimization problems are hard to solve. The notion of problem hardness is captured by the theory of computational complexity [49,74] and for many important problems it is well known that they are \mathcal{NP} -hard, that is, the time needed to solve an instance in the worst case grows exponentially with instance size. Often, approximate algorithms are the only feasible way to obtain near optimal solutions at relatively low computational cost.

Most approximate algorithms are either *construction* algorithms or *local search* algorithms.¹ These two types of methods are significantly different, because construction algorithms work on partial solutions trying to extend these in the best possible way to complete problem solutions, while local search methods move in the search space of complete solutions.

¹Other approximate methods are also conceivable. For example, when stopping exact methods, like Branch & Bound, before completion [4,58] (e.g., after some given time bound, or when some guarantee on the solution quality is obtained through the use of lower and upper bounds), we can convert exact algorithms into approximate ones.

2.1 Construction Algorithms

Construction algorithms build solutions to a problem under consideration in an incremental way starting with an empty initial solution and iteratively adding appropriate solution components without backtracking until a complete solution is obtained. In the simplest case, solution components are added in random order. Often better results are obtained if a heuristic estimate of the myopic benefit of adding solution components is taken into account. *Greedy construction heuristics* add at each step a solution component which achieves the maximal myopic benefit as measured by some heuristic information. An algorithmic outline of a greedy construction heuristic is given in Figure 9.1. The function `GreedyComponent` returns the solution component e with the best heuristic estimate. Solutions returned by greedy algorithms are typically of better quality than randomly generated solutions. Yet, a disadvantage of greedy construction heuristics is that they can generate only a very limited number of different solutions. Additionally, greedy decisions in early stages of the construction process strongly constrain the available possibilities at later stages, often causing very poor moves in the final phases of the solution construction.

As an example, consider a greedy construction heuristic for the traveling salesman problem. In the TSP we are given a complete weighted graph $G = (\mathcal{N}, \mathcal{A})$ with \mathcal{N} being the set of vertices, representing the cities, and \mathcal{A} the set of edges fully connecting the vertices \mathcal{N} . Each edge is assigned a value d_{ij} , which is the length of edge $(i, j) \in \mathcal{A}$. The TSP is the problem of finding a minimal length Hamiltonian circuit of the graph, where an Hamiltonian circuit is a closed tour visiting exactly once each of the $n = |\mathcal{N}|$ vertices of G . For symmetric TSPs, the distances between the cities are independent of the direction of traversing the edges, that is, $d_{ij} = d_{ji}$ for every pair of vertices. In the more general asymmetric TSP (ATSP) at least for one pair of vertices i, j we have $d_{ij} \neq d_{ji}$.

A simple rule of thumb to build a tour is to start from some initial city and to always choose to go to the closest still unvisited city before returning to the start city. This algorithm is known as the *nearest neighbor* tour construction heuristic.

Figure 9.2 shows a tour returned by the nearest neighbor heuristic on TSP instance `att532`, taken from `TSPLIB`,² with 532 cities in the US. Noteworthy in this example is that there are a few very long links in the tour, leading to a strongly suboptimal solution.

```

procedure Greedy Construction Heuristic
   $s_p = \text{empty solution}$ 
  while  $s_p$  not_a_complete_solution do
     $e = \text{GreedyComponent}(s_p)$ 
     $s_p = s_p \otimes e$ 
  end
  return  $s_p$ 
end Greedy Construction Heuristic

```

Figure 9.1. Algorithmic skeleton of a greedy construction heuristic. The addition of component e to a partial solution s_p is denoted by the operator \otimes .

²TSPLIB is a benchmark library for the TSP and related problems and is accessible via <http://www.iwr.uni-heidelberg.de/iwr/comopt/software/TSPLIB95>.

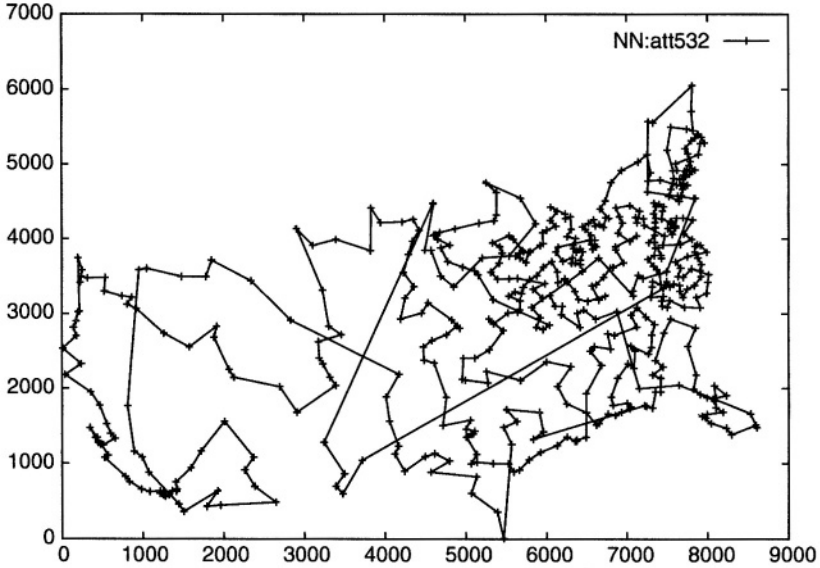


Figure 9.2. Tour returned by the nearest neighbor heuristic on TSP instance att532 from TSPLIB.

In fact, construction algorithms are typically the fastest approximate methods, but the solutions they generate often are not of a very high quality and they are not guaranteed to be optimal with respect to small changes; the results produced by constructive heuristics can often be improved by local search algorithms.

2.2 Local Search

Local search algorithms start from a complete initial solution and try to find a better solution in an appropriately defined *neighborhood* of the current solution. In its most basic version, known as *iterative improvement*, the algorithm searches the neighborhood for an improving solution. If such a solution is found, it replaces the current solution and the local search continues. These steps are repeated until no improving neighbor solution can be found in the neighborhood of the current solution and the algorithm ends in a *local optimum*. An outline of an iterative improvement algorithm is given in Figure 9.3. The procedure `Improve` returns a better neighbor solution if one exists, otherwise it returns the current solution, in which case the algorithm stops.

The choice of an appropriate neighborhood structure is crucial for the performance of the local search algorithm and has to be done in a problem specific way. The neighborhood structure defines the set of solutions that can be reached from s in one single step of a local search algorithm. An example neighborhood for the TSP is the k -change neighborhood in which neighbor solutions differ by at most k edges. Figure 9.4 shows an example of a 2-change neighborhood. The 2-change algorithm systematically tests whether the current tour can be improved by replacing two edges. To fully specify a local search algorithm it is necessary to designate a particular neighborhood examination scheme that defines how the neighborhood is searched and which neighbor

```

procedure IterativeImprovement ( $s \in \mathcal{S}$ )
   $s' = \text{Improve}(s)$ 
  while  $s' \neq s$  do
     $s = s'$ 
     $s' = \text{Improve}(s)$ 
  end
  return  $s$ 
end IterativeImprovement

```

Figure 9.3. Algorithmic skeleton of iterative improvement.




Figure 9.4. Schematic illustration of a 2-change move. The proposed move reduces the total tour length if we consider the Euclidean distance between the points.

solution replaces the current one. In the case of iterative improvement algorithms, this rule is called the *pivoting rule* [93] and examples are the *best-improvement* rule, which chooses the neighbor solution giving the largest improvement of the objective function, and the *first-improvement* rule, which uses the first improved solution found in the neighborhood to replace the current one. A common problem with local search algorithms is that they easily get trapped in local minima and that the result strongly depends on the initial solution.

3 THE ACO METAHEURISTIC

Artificial ants used in ACO are stochastic solution construction procedures that probabilistically build a solution by iteratively adding solution components to partial solutions by taking into account (i) heuristic information on the problem instance being solved, if available, and (ii) (artificial) pheromone trails which change dynamically at run-time to reflect the agents' acquired search experience.

A stochastic component in ACO allows the ants to build a wide variety of different solutions and hence explore a much larger number of solutions than greedy heuristics. At the same time, the use of heuristic information, which is readily available for many problems, can guide the ants towards the most promising solutions. More important, the ants' search experience can be used to influence, in a way reminiscent of reinforcement learning [89], the solution construction in future iterations of the algorithm. Additionally, the use of a colony of ants can give the algorithm increased robustness and in many ACO applications the collective interaction of a population of agents is needed to efficiently solve a problem.

The domain of application of ACO algorithms is vast. In principle, ACO can be applied to any discrete optimization problem for which some solution construction

mechanism can be conceived. In the following of this section, we first define a generic problem representation which the ants in ACO exploit to construct solutions, then we detail the ants' behavior while constructing solutions, and finally we define the ACO metaheuristic.

3.1 Problem Representation

Let us consider the minimization problem³ (\mathcal{S}, f, Ω) , where \mathcal{S} is the *set of candidate solutions*, f is the *objective function* which assigns to each candidate solution $s \in \mathcal{S}$ an objective function (cost) value $f(s, t)$,⁴ and Ω is a set of constraints. The goal is to find a *globally optimal* solution $s_{opt} \in \mathcal{S}$, that is, a minimum cost solution that satisfies the constraints Ω .

The problem representation of a combinatorial optimization problem (\mathcal{S}, f, Ω) which is exploited by the ants can be characterized as follows:

- A finite set $\mathcal{C} = \{c_1, c_2, \dots, c_{N_C}\}$ of *components* is given.
- The *states* of the problem are defined in terms of sequences $x = \langle c_i, c_j, \dots, c_k, \dots \rangle$ over the elements of \mathcal{C} . The set of all possible sequences is denoted by \mathcal{X} . The length of a sequence x , that is, the number of components in the sequence, is expressed by $|x|$. The maximum length of a sequence is bounded by a positive constant $n < +\infty$.
- The set of (candidate) solutions \mathcal{S} is a subset of \mathcal{X} (i.e., $\mathcal{S} \subseteq \mathcal{X}$).
- The finite set of *constraints* Ω defines the set of feasible states $\tilde{\mathcal{X}}$, with $\tilde{\mathcal{X}} \subseteq \mathcal{X}$.
- A non-empty set \mathcal{S}^* of feasible solutions is given, with $\mathcal{S}^* \subseteq \tilde{\mathcal{X}}$ and $\mathcal{S}^* \subseteq \mathcal{S}$.
- A *cost* $f(s, t)$ is associated to each candidate solution $s \in \mathcal{S}$.
- In some cases a cost, or the estimate of a cost, $J(x_i, t)$ can be associated to states other than solutions. If x_j can be obtained by adding solution components to a state x_i then $J(x_i, t) \leq J(x_j, t)$. Note that $J(s, t) \equiv f(s, t)$.

Given this representation, artificial ants build solutions by moving on the construction graph $G = (\mathcal{C}, \mathcal{L})$, where the vertices are the components \mathcal{C} and the set \mathcal{L} fully connects \mathcal{C} (elements of \mathcal{L} are called *connections*). The problem constraints Ω are implemented in the policy followed by the artificial ants build solutions, as explained in the next section. The choice of implementing the constraints in the construction policy of the artificial ants allows a certain degree of flexibility. In fact, depending on the combinatorial optimization problem considered, it may be more reasonable to implement constraints in a hard way allowing ants to build only feasible solutions, or in a soft way, in which case ants can build infeasible solutions (that is, candidate solutions in $\mathcal{S} \setminus \mathcal{S}^*$) that will be penalized, depending on their degree of infeasibility.

3.2 Ant's Behavior

Ants can be characterized as stochastic construction procedures which build solutions moving on the construction graph $G = (\mathcal{C}, \mathcal{L})$. Ants do not move arbitrarily on G ,

³The adaptation to a maximization problem is straightforward.

⁴The parameter t indicates that the objective function can be time dependent, as it is the case in applications to dynamic problems.

but rather follow a construction policy which is a function of the problem constraints Ω . In general, ants try to build feasible solutions, but, if necessary, they can generate infeasible solutions. Components $c_i \in \mathcal{C}$ and connections $l_{ij} \in \mathcal{L}$ can have associated a *pheromone trail* τ (τ_i if associated to components, τ_{ij} if associated to connections) encoding a long-term memory about the whole ant search process that is updated by the ants themselves, and a *heuristic value* η (η_i and η_{ij} , respectively) representing a priori information about the problem instance definition or run-time information provided by a source different from the ants. In many cases η is the cost, or an estimate of the cost, of extending the current state. These values are used by the ants' heuristic rule to make probabilistic decisions on how to move on the graph.

More precisely, each ant k of the colony has the following properties:

- It exploits the graph $G = (\mathcal{C}, \mathcal{L})$ to search for feasible solutions s of minimum cost. That is, solutions s such that $\hat{f}_s = \min_s f(s, t)$.
- It has a memory \mathcal{M}^k that it uses to store information about the path it followed so far. Memory can be used (i) to build feasible solutions (i.e., to implement constraints Ω), (ii) to evaluate the solution found, and (iii) to retrace the path backward to deposit pheromone.
- It can be assigned a *start state* x_s^k and one or more *termination conditions* e^k . Usually, the start state is expressed either as a unit length sequence (that is, a single component sequence), or an empty sequence.
- When in state $x_r = \langle x_{r-1}, i \rangle$ it tries to move to any vertex j in its *feasible neighborhood* \mathcal{N}_i^k , that is, to a state $\langle x_r, j \rangle \in \tilde{\mathcal{X}}$. If this is not possible, then the ant might be allowed to move to a vertex j in its infeasible neighborhood \mathcal{IN}_i^k , generating in this way an infeasible state (that is, a state $\langle x_r, j \rangle \in \mathcal{X} \setminus \tilde{\mathcal{X}}$).
- It selects the move by applying a probabilistic decision rule. Its probabilistic decision rule is a function of (i) locally available pheromone trails and heuristic values, (ii) the ant's private memory storing its past history, and (iii) the problem constraints.
- The construction procedure of ant k stops when at least one of the termination conditions e^k is satisfied. Examples of termination conditions are when a solution is completed, or when, if building infeasible solutions is not allowed, there are no feasible states reachable from the ant current state.
- When adding a component c_j to the current solution it can update the pheromone trail associated to it or to the corresponding connection. This is called *online step-by-step pheromone update*.
- Once built a solution, it can retrace the same path backward and update the pheromone trails of the used components or connections. This is called *online delayed pheromone update*.

It is important to note that ants move concurrently and independently and that each ant is complex enough to find a (probably poor) solution to the problem under consideration. Typically, good quality solutions emerge as the result of the collective interaction among the ants which is obtained via indirect communication mediated by the information ants read/write in the variables storing pheromone trail values. In a way, this is a distributed learning process in which the single agents, the ants, are not

adaptive themselves but, on the contrary, they adaptively modify the way the problem is represented and perceived by other ants.

3.3 The Metaheuristic

Informally, the behavior of ants in an ACO algorithm can be summarized as follows. A colony of ants concurrently and asynchronously move through adjacent states of the problem by building paths on G . They move by applying a stochastic local decision policy that makes use of pheromone trails and heuristic information. By moving, ants incrementally build solutions to the optimization problem. Once an ant has built a solution, or while the solution is being built, the ant evaluates the (partial) solution and deposits pheromone trails on the components or connections it used. This pheromone information will direct the search of future ants.

Besides ants' activity, an ACO algorithm includes two additional procedures: *pheromone trail evaporation* and *daemon actions* (this last component being optional). Pheromone evaporation is the process by means of which the pheromone deposited by previous ants decreases over time. From a practical point of view, pheromone evaporation is needed to avoid a too rapid convergence of the algorithm towards a sub-optimal region. It implements a useful form of *forgetting*, favoring the exploration of new areas of the search space. Daemon actions can be used to implement centralized actions which cannot be performed by single ants. Examples are the activation of a local optimization procedure, or the collection of global information that can be used to decide whether it is useful or not to deposit additional pheromone to bias the search process from a non-local perspective. As a practical example, the daemon can observe the path found by each ant in the colony and choose to deposit extra pheromone on the components used by the ant that built the best solution. Pheromone updates performed by the daemon are called *off-line pheromone updates*.

In Figure 9.5 the ACO metaheuristic behavior is described in pseudo-code. The main procedure of the ACO metaheuristic manages, via the **ScheduleActivities** construct, the scheduling of the three above discussed components of ACO algorithms: (i) management of ants' activity, (ii) pheromone evaporation, and (iii) daemon actions. The **ScheduleActivities** construct does not specify how these three activities are scheduled and synchronized. In other words, it does not say whether they should be executed in a completely parallel and independent way, or if some kind of synchronization

```

procedure ACO metaheuristic
  ScheduleActivities
    ManageAntsActivity()
    EvaporatePheromone()
    DaemonActions() {Optional}
  end ScheduleActivities
end ACO metaheuristic

```

Figure 9.5. The ACO metaheuristic in pseudo-code. Comments are enclosed in braces. The procedure `DaemonActions()` is optional and refers to centralized actions executed by a daemon possessing global knowledge.

among them is necessary. The designer is therefore free to specify the way these three procedures should interact.

4 HISTORY OF ACO ALGORITHMS

The first ACO algorithm proposed was Ant System (AS). AS was applied to some rather small instances of the traveling salesman problem (TSP) with up to 75 cities. It was able to reach the performance of other general-purpose heuristics like evolutionary computation [30,38]. Despite these initial encouraging results, AS did not prove to be competitive with state-of-the-art algorithms specifically designed for the TSP when attacking large instances. Therefore, a substantial amount of recent research has focused on ACO algorithms which show better performance than AS when applied, for example, to the TSP. In the following of this section we first briefly introduce the biological metaphor on which AS and ACO are inspired, and then we present a brief history of the developments that have led from the original AS to the most recent ACO algorithms. In fact, these more recent algorithms are direct extensions of AS which add advanced features to improve the algorithm performance.

4.1 Biological Analogy

In many ant species, individual ants may deposit a pheromone (a particular chemical that ants can smell) on the ground while walking [23,52]. By depositing pheromone they create a trail that is used, e.g., to mark the path from the nest to food sources and back. In fact, by sensing pheromone trails foragers can follow the path to food discovered by other ants. Also, they are capable of exploiting pheromone trails to choose the shortest among the available paths leading to the food.

Deneubourg and colleagues [23,52] used a double bridge connecting a nest of ants and a food source to study pheromone trail laying and following behavior in controlled experimental conditions.⁵ They ran a number of experiments in which they varied the ratio between the length of the two branches of the bridge. The most interesting, for our purposes, of these experiments is the one in which one branch was longer than the other. In this experiment, at the start the ants were left free to move between the nest and the food source and the percentage of ants that chose one or the other of the two branches was observed over time. The outcome was that, although in the initial phase random oscillations could occur, in most experiments all the ants ended up using the shorter branch.

This result can be explained as follows. When a trial starts there is no pheromone on the two branches. Hence, the ants do not have a preference and they select with the same probability either of the two branches. Therefore, it can be expected that, on average, half of the ants choose the short branch and the other half the long branch, although stochastic oscillations may occasionally favor one branch over the other. However, because one branch is shorter than the other, the ants choosing the short branch are the first to reach the food and to start their travel back to the nest.⁶ But then, when they must

⁵The experiment described was originally executed using a laboratory colony of Argentine ants (*Iridomyrmex humilis*). It is known that these ants deposit pheromone both when leaving and when returning to the nest [52].

⁶In the ACO literature this is often called *differential path length* effect.

make a decision between the short and the long branch, the higher level of pheromone on the short branch biases their decision in its favor.⁷ Therefore, pheromone starts to accumulate faster on the short branch which will eventually be used by the great majority of the ants.

It should be clear by now how real ants have inspired AS and later algorithms: the double bridge was substituted by a graph, and pheromone trails by artificial pheromone trails. Also, because we wanted artificial ants to solve problems more complicated than those solved by real ants, we gave artificial ants some extra capacities, like a memory (used to implement constraints and to allow the ants to retrace their path back to the nest without errors) and the capacity for depositing a quantity of pheromone proportional to the quality of the solution produced (a similar behavior is observed also in some real ants species in which the quantity of pheromone deposited while returning to the nest from a food source is proportional to the quality of the food source found [3]).

In the next section we will see how, starting from AS, new algorithms have been proposed that, although retaining some of the original biological inspiration, are less and less biologically inspired and more and more motivated by the need of making ACO algorithms competitive with or indeed better than other state-of-the-art algorithms. Nevertheless, many aspects of the original Ant System remain: the need for a colony, the role of autocatalysis, the cooperative behavior mediated by artificial pheromone trails, the probabilistic construction of solutions biased by artificial pheromone trails and local heuristic information, the pheromone updating guided by solution quality, and the evaporation of pheromone trail, are present in all ACO algorithms. It is interesting to note that there is one well known algorithm that, although making use in some way of the ant foraging metaphor, cannot be considered an instance of the Ant Colony Optimization metaheuristic. This is HAS-QAP, proposed in [48], where pheromone trails are not used to guide the solution construction phase; on the contrary, they are used to guide modifications of complete solutions in a local search style. This algorithm belongs nevertheless to *ant algorithms*, a new class of algorithms inspired by a number of different behaviors of social insects. Ant algorithms are receiving increasing attention in the scientific community (see, e.g., [8,9,11,31]) as a promising novel approach to distributed control and optimization.

4.2 Historical Development

As we said, AS was the first example of an ACO algorithm to be proposed in the literature. In fact, AS was originally a set of three algorithms called *ant-cycle*, *ant-density*, and *ant-quantity*. These three algorithms were proposed in Dorigo's doctoral dissertation [30] and first appeared in a technical report [36,37] that was published a few years later in the IEEE Transactions on Systems, Man, and Cybernetics [38]. Another early publication is [17].

While in *ant-density* and *ant-quantity* the ants updated the pheromone directly after a move from a city to an adjacent one, in *ant-cycle* the pheromone update was only done after all the ants had constructed the tours and the amount of pheromone deposited by each ant was set to be a function of the tour quality. Because *ant-cycle* performed better than the other two variants, it was later called simply Ant System (and in fact, it

⁷ A process like this, in which a decision taken at time t increases the probability of making the same decision at time $T > t$ is said to be an *autocatalytic* process. Autocatalytic processes exploit *positive feedback*.

is the algorithm that we will present in the following subsection), while the other two algorithms were no longer studied.

The major merit of AS, whose computational results were promising but not competitive with other more established approaches, was to stimulate a number of researchers, mostly in Europe, to develop extensions and improvements of its basic ideas so as to produce better performing, and often state-of-the-art, algorithms. It is following the successes of this collective undertaking that recently Dorigo and Di Caro [32] made the synthesis effort that led to the definition of the ACO metaheuristic presented in this chapter (see also [33]). In other words, the ACO metaheuristic was defined *a posteriori* with the goal of providing a common characterization of a new class of algorithms and a reference framework for the design of new instances of ACO algorithms.

4.2.1 The First ACO Algorithm: Ant System and the TSP

The traveling salesman problem (TSP) is a paradigmatic \mathcal{NP} -hard combinatorial optimization problem which has attracted an enormous amount of research effort [57,60,76]. The TSP is a very important problem also in the context of Ant Colony Optimization because it is the problem to which the original AS was first applied [30,36,38], and it has later often been used as a benchmark to test new ideas and algorithmic variants.

In AS each ant is initially put on a randomly chosen city and has a memory which stores the partial solution it has constructed so far (initially the memory contains only the start city). Starting from its start city, an ant iteratively moves from city to city. When at a city i , an ant k chooses to go to an as yet unvisited city j with a probability given by

$$p_{ij}^k(t) = \frac{[\tau_{ij}(t)]^\alpha \cdot [\eta_{ij}]^\beta}{\sum_{l \in \mathcal{N}_i^k} [\tau_{il}(t)]^\alpha \cdot [\eta_{il}]^\beta} \quad \text{if } j \in \mathcal{N}_i^k \quad (9.1)$$

where $\eta_{ij} = 1/d_{ij}$ is *a priori* available heuristic information, α and β are two parameters which determine the relative influence of pheromone trail and heuristic information, and \mathcal{N}_i^k is the feasible neighborhood of ant k , that is, the set of cities which ant k has not yet visited. Parameters α and β have the following influence on the algorithm behavior. If $\alpha = 0$, the selection probabilities are proportional to $[\eta_{ij}]^\beta$ and the closest cities will more likely be selected: in this case AS corresponds to a classical stochastic greedy algorithm (with multiple starting points since ants are initially randomly distributed on the cities). If $\beta = 0$, only pheromone amplification is at work: this will lead to the rapid emergence of a *stagnation* situation with the corresponding generation of tours which, in general, are strongly suboptimal [30]. (Search stagnation is defined in [38] as the situation where all the ants follow the same path and construct the same solution.)

The solution construction ends after each ant has completed a tour, that is, after each ant has constructed a sequence of length n . Next, the pheromone trails are updated. In AS this is done by first lowering the pheromone trails by a constant factor (this is pheromone evaporation) and then allowing each ant to deposit pheromone on the edges that belong to its tour:

$$\forall(i, j) \quad \tau_{ij}(t + 1) = (1 - \rho) \cdot \tau_{ij}(t) + \sum_{k=1}^m \Delta \tau_{ij}^k(t) \quad (9.2)$$

where $0 < \rho \leq 1$ is the pheromone trail evaporation rate and m is the number of ants. The parameter ρ is used to avoid unlimited accumulation of the pheromone trails

and enables the algorithm to “forget” previous “bad” decisions. On edges which are not chosen by the ants, the associated pheromone strength will decrease exponentially with the number of iterations. $\Delta\tau_{ij}^k(t)$ is the amount of pheromone ant k deposits on the edges; it is defined as

$$\Delta\tau_{ij}^k(t) = \begin{cases} 1/L^k(t) & \text{if edge } (i, j) \text{ is used by ant } k \\ 0 & \text{otherwise} \end{cases} \quad (9.3)$$

where $L^k(t)$ is the length of the k th ant’s tour. By Equation 9.3, the shorter the ant’s tour is, the more pheromone is received by edges belonging to the tour.⁸ In general, edges which are used by many ants and which are contained in shorter tours will receive more pheromone and therefore are also more likely to be chosen in future iterations of the algorithm.

4.2.2 Ant System and its Extensions

As previously stated, AS was not competitive with state-of-the-art algorithms for TSP. Researchers then started to extend it to try to improve its performance.

A first improvement, called the *elitist strategy*, was introduced in [30,38]. It consists in giving the best tour since the start of the algorithm (called T^{gb} , where gb stands for global-best) a strong additional weight. In practice, each time the pheromone trails are updated, those belonging to the edges of the global-best tour get an additional amount of pheromone. For these edges Equation 9.3 becomes:

$$\Delta\tau_{ij}^{gb}(t) = \begin{cases} e/L^{gb}(t) & \text{if edge } (i, j) \in T^{gb} \\ 0 & \text{otherwise} \end{cases} \quad (9.4)$$

The edges of T^{gb} are therefore reinforced with a quantity of pheromone given by e/L^{gb} , where L^{gb} is the length of T^{gb} and e is a positive integer. Note that this type of pheromone update is a first example of daemon action as described in Section 3.3.

Other improvements were the rank-based version of Ant System (AS_{rank}), $MMAS$ – $MMAS$ Ant System ($MMAS$), and Ant Colony System (ACS). AS_{rank} [14] is in a sense an extension of the elitist strategy: it sorts the ants according to the lengths of the tours they generated and, after each tour construction phase, only the $(w - 1)$ best ants and the global-best ant are allowed to deposit pheromone. The r th best ant of the colony contributes to the pheromone update with a weight given by $\max\{0, w - r\}$ while the global-best tour reinforces the pheromone trails with weight w , Equation 9.2 becomes therefore:

$$\forall(i, j) \quad \tau_{ij}(t + 1) = (1 - \rho) \cdot \tau_{ij}(t) + \sum_{r=1}^{w-1} (w - r) \cdot \Delta\tau_{ij}^r(t) + w \cdot \Delta\tau_{ij}^{gb}(t) \quad (9.5)$$

where $\Delta\tau_{ij}^r(t) = 1/L^r(t)$ and $\Delta\tau_{ij}^{gb}(t) = 1/L^{gb}$.

ACS [34,35,44] improves over AS by increasing the importance of exploitation of information collected by previous ants with respect to exploration of the search

⁸Note that when applied to symmetric TSPs the edges are considered to be bidirectional and edges (i, j) and (j, i) are both updated. This is different for the ATSP, where edges are directed; an ant crossing edge (i, j) will update only this edge and not the edge (j, i) .

space.⁹ This is achieved via two mechanisms. First, a strong elitist strategy is used to update pheromone trails. Second, ants choose the next city to move to using a so-called *pseudo-random proportional* rule [35]: with probability q_0 , $0 \leq q_0 \leq 1$, they move to the city j for which the product between pheromone trail and heuristic information is maximum, that is, $j = \arg \max_{j \in \mathcal{N}_i^k} \{\tau_{ij}(t) \cdot \eta_{ij}^\beta\}$, while with probability $1 - q_0$ they operate a biased exploration in which the probability $p_{ij}^k(t)$ is the same as in AS (see Equation 9.1). The value q_0 is a parameter: when it is set to a value close to 1, as it is the case in most ACS applications, exploitation is favored over exploration. Obviously, when $q_0 = 0$ the probabilistic decision rule becomes the same as in AS.

As mentioned earlier, pheromone updates are performed using a strong elitist strategy: only the ant that has produced the best solution is allowed to update pheromone trails, according to a pheromone trail update rule similar to that used in AS:

$$\tau_{ij}(t+1) = (1 - \rho) \cdot \tau_{ij}(t) + \rho \cdot \Delta \tau_{ij}^{\text{best}}(t) \quad (9.6)$$

The best ant can be the *iteration-best* ant, that is, the best in the current iteration, or the *global-best* ant, that is, the ant that made the best tour from the start of the trial.

Finally, ACS differs from previous ACO algorithms also because ants update the pheromone trails while building solutions (as in *ant-quantity* and in *ant-density*). In practice, ACS ants “eat” some of the pheromone trail on the edges they visit. This has the effect of decreasing the probability that the same path is used by all the ants (i.e., it favors exploration, counterbalancing this way the other two above-mentioned modifications that strongly favor exploitation of the collected knowledge about the problem). ACS has been improved also by the addition of local search routines that take the solution generated by ants to their local optimum just before the pheromone update.

MMAS [84,87,88] introduces upper and lower bounds to the values of the pheromone trails, as well as a different initialization of their values. In practice, in *MMAS* the allowed range of the pheromone trail strength is limited to the interval $[\tau_{\min}, \tau_{\max}]$, that is, $\tau_{\min} \leq \tau_{ij} \leq \tau_{\max} \forall \tau_{ij}$, and the pheromone trails are initialized to the upper trail limit, which causes a higher exploration at the start of the algorithm. Also, as in ACS, in *MMAS* only the best ant (the global-best or the iteration-best ant) is allowed to add pheromone after each algorithm iteration. In fact, in *MMAS* the iteration-best ant and the global-best ant can be used alternately in the pheromone update. Computational results have shown that best results are obtained when pheromone updates are performed using the global-best solution with increasing frequency during the algorithm execution. Similarly to ACS, also *MMAS* often exploits local search to improve its performance.

4.2.3 Applications to Dynamic Problems

The application of ACO algorithms to dynamic problems, that is, problems whose characteristics change while being solved, is the most recent major development in the field. The first such application [79] concerned routing in circuit-switched networks

⁹ACS was an offspring of Ant-Q [43], an algorithm intended to create a link between reinforcement learning [89] and Ant Colony Optimization. Computational experiments have shown that some aspects of Ant-Q, in particular the pheromone update rule, could be strongly simplified without affecting performance. It is for this reason that Ant-Q was abandoned in favor of the simpler and equally good ACS.

(e.g., classical telephone networks). The proposed algorithm, called ABC, was demonstrated on a simulated version of the British Telecom network. The main merit of ABC was to stimulate the interest of ACO researchers in dynamic problems. In fact, only rather limited comparisons were made between ABC and state-of-the-art algorithms, so that it is not possible to judge on the quality of the results obtained.

A very successful application of ACO to dynamic problems is the AntNet algorithm, proposed by Di Caro and Dorigo [24–26,28] and discussed in Section 5. AntNet was applied to routing in packet-switched networks (e.g., the Internet). It contains a number of innovations with respect to AS and it has been shown experimentally to outperform a whole set of state-of-the-art algorithms on numerous benchmark problems.

5 EXAMPLES OF APPLICATIONS

The versatility and the practical use of the ACO metaheuristic for the solution of combinatorial optimization problems is best illustrated via example applications to a number of different problems.

The ACO application to the TSP has already been presented in the previous section. Here, we additionally discuss applications to three \mathcal{NP} -hard optimization problems, the single machine total weighted tardiness problem (SMTWTP), the generalized assignment problems (GAP), and the set covering problem (SCP). We have chosen these problems to make the application examples as comprehensive as possible with respect to different ways of representing solutions. While the TSP and the SMTWTP are permutation problems, that is, solutions are represented as permutations of solution components, solutions in the GAP are assignments of tasks to agents and in the SCP a solution is represented as a subset of the available solution components.

Applications of ACO to dynamic problems focus mainly on routing in data networks. As an example, in the following we present the AntNet algorithm [26].

Example 9.1. *The single machine total weighted tardiness scheduling problem (SMTWTP)*

In the SMTWTP n jobs have to be processed sequentially without interruption on a single machine. Each job has an associated processing time p_j , a weight w_j , and a due date d_j and all jobs are available for processing at time zero. The tardiness of job j is defined as $T_j = \max\{0, C_j - d_j\}$, where C_j is its completion time in the current job sequence. The goal in the SMTWTP is to find a job sequence which minimizes the sum of the weighted tardiness given by $\sum_{i=1}^n w_i \cdot T_i$.

For the ACO application to the SMTWTP, the set of components \mathcal{C} is the set of all jobs. As in the TSP case, the states of the problem are all possible partial sequences. In the SMTWTP case we do not have explicit costs associated with the connections because the objective function contribution of each job depends on the partial solution constructed so far.

The SMTWTP was attacked in [22] using ACS (ACS-SMTWTP). In ACS-SMTWTP, each ant starts with an empty sequence and then iteratively appends an unscheduled job to the partial sequence constructed so far. Each ant chooses the next job using the *pseudo-random-proportional* action choice rule, where at each step the feasible neighborhood \mathcal{N}_i^k of ant k is the list of as yet unscheduled jobs. Pheromone trails are defined as follows: τ_{ij} refers to the desirability of scheduling job j at position i . This definition of the pheromone trails is, in fact, used in most ACO application

to scheduling problems [2,22,67,82]. Concerning the heuristic information, in [22] the use of three priority rules allowed to define three different types of heuristic information for the SMTWTP. The investigated priority rules were: (i) the earliest due date rule, which puts the jobs in non-decreasing order of the due dates d_j , (ii) the modified due date rule which puts the jobs in non-decreasing order of the modified due dates given by $mdd_j = \max\{C + p_j, d_j\}$ [2], where C is the sum of the processing times of the already sequenced jobs, and (iii) the apparent urgency rule which puts the jobs in non-decreasing order of the apparent urgency [72], given by $au_j = (w_j/p_j) \cdot \exp(-(\max\{d_j - C_j, 0\})/k\bar{p})$, where k is a parameter of the priority rule. In each case, the heuristic information was defined as $\eta_{ij} = 1/h_j$, where h_j is either d_j , mdd_j , or au_j , depending on the priority rule used.

The global and the local pheromone updates are carried out as in the standard ACS described in Section 4.2, where in the global pheromone update, T^{gb} is the total weighted tardiness of the global best solution.

In [22], ACS-SMTWTP was combined with a powerful local search algorithm. The final ACS algorithm was tested on a benchmark set available from ORLIB at <http://www.ms.ic.ac.uk/info.html>. Within the computation time limits given,¹⁰ ACS reached a very good performance and could find in each single run the optimal or best known solutions on all instances of the benchmark set. For more details on the computational results we refer to [22].

Example 9.2. *The generalized assignment problem (GAP)*

In the GAP a set of tasks $\mathcal{I}, i = 1, \dots, n$, has to be assigned to a set of agents $\mathcal{J}, j = 1, \dots, m$. Each agent j has only a limited capacity a_j and each task i consumes, when assigned to agent j , a quantity b_{ij} of the agent’s capacity. Also, the cost d_{ij} of assigning task i to agent j is given. The objective then is to find a feasible task assignment with minimal cost.

Let y_{ij} be one if task i is assigned to agent j and zero otherwise. Then the GAP can formally be defined as

$$\min f(y) = \sum_{i=1}^n \sum_{j=1}^m d_{ij} \cdot y_{ij} \tag{9.7}$$

subject to

$$\sum_{i=1}^n b_{ij} \cdot y_{ij} \leq a_j \quad j = 1, \dots, m \tag{9.8}$$

$$\sum_{j=1}^m y_{ij} = 1 \quad i = 1, \dots, n \tag{9.9}$$

$$y_{ij} \in \{0, 1\} \quad i = 1 \dots, n; \quad j = 1, \dots, m \tag{9.10}$$

The constraints 9.8 implement the limited resource capacity of the agents, while constraints 9.9 and 9.10 impose that each task is assigned to exactly one agent and that a task cannot be split among several agents.

¹⁰The maximum time for the largest instances was 20min on a 450 MHz Pentium III PC with 256 MB RAM. Programs were written in C++ and the PC was run under Red Hat Linux 6.1.

The GAP can easily be cast into the framework of the ACO metaheuristic. The problem can be represented by a graph in which the set of components comprises the set of tasks and agents, that is $\mathcal{C} = \mathcal{I} \cup \mathcal{J}$ and the set of connections fully connect the graph. Each assignment, which consists of n couplings (i, j) of tasks and agents, corresponds to an ant's walk on this graph. Such a walk has to observe the constraints 9.9 and 9.10 to obtain a valid assignment. One particular way of generating such an assignment is by an ant's walk which iteratively switches from task vertices (vertices in the set \mathcal{J}) to agent vertices (vertices in the set \mathcal{I}) without repeating any task vertex but possibly using the same agent vertex several times (that is, several tasks may be assigned to the same agent).

At each step of the construction process, an ant has to make one of the following two basic decisions: (i) it has to decide which task to assign next and (ii) it has to decide to which agent a chosen task should be assigned. Pheromone trail and heuristic information can be associated with both steps. With respect to the first step the pheromone information can be used to learn an appropriate assignment order of the tasks, that is, τ_{ij} gives the desirability of assigning task j directly after task i , while the pheromone information in the second step is associated with the desirability of assigning a task to a specific agent.

For simplicity let us consider an approach in which the tasks are assigned in a random order. Then, at each step a task has to be assigned to an agent. Intuitively, it is better to assign tasks to agents such that small assignment costs are incurred and that the agent needs only a relatively small amount of its available resource to perform the task. Hence, one possible heuristic information is $\eta_{ij} = a_j/d_{ij}b_{ij}$ and a probabilistic selection of the assignments can follow the AS probabilistic rule (Equation 9.1) or the pseudo-random proportional rule of ACS. Yet, a complication in the construction process is that the GAP involves resource capacity constraints and, in fact, for the GAP no guarantee is given that an ant will construct a feasible solution which obeys the resource constraints given by Equation 9.8. In fact, to have a bias towards generating feasible solutions, the resource constraints should be taken into account in the definition of \mathcal{N}_i^k , the feasible neighborhood of ant k . For the GAP, we define \mathcal{N}_i^k to consist of all those agents to which the task i can be assigned without violating the agents' resource capacity. If no agent can meet the task's resource requirement, we are forced to build an infeasible solution and in this case we can simply choose \mathcal{N}_i^k as the set of all agents. Infeasibilities can then be handled, for example, by assigning penalties proportional to the amount of resource violations.

A first application of *MAX-MIN* Ant System (*MMAS*) to the GAP was presented in [75]. The approach shows some particularities, like the use of a single ant and the lack of any heuristic information. The infeasibility of solutions is only treated in the pheromone update: the amount of pheromone deposited by an ant is set to a high value if the solution it generated is feasible, to a low value if it is infeasible. These values are constants independent of the solution quality. Additionally, *MMAS* was coupled with a local search based on tabu search and ejection chain elements [51] and it obtained very good performance on benchmark instances available at ORLIB.

Example 9.3. *The set covering problem (SCP)*

In the set covering problem (SCP) we are given a $m \times n$ matrix $A = (a_{ji})$ in which all the matrix elements are either zero or one. Additionally, each column is given a non-negative cost b_i . We say that a column i covers a row j if $a_{ji} = 1$. The goal in the

SCP is to choose a subset of the columns of minimal weight that covers every row. Let \mathcal{I} denote a subset of the columns and y_i be a binary variable which is one, if $i \in \mathcal{I}$, and zero otherwise. The SCP can be defined formally as follows.

$$\min f(y) = \sum_{i=1}^n b_i \cdot y_i \tag{9.11}$$

subject to

$$\sum_{i=1}^n a_{ji} \cdot y_i \geq 1 \quad j = 1, \dots, m \tag{9.12}$$

$$y_i \in \{0, 1\} \quad i = 1, \dots, n \tag{9.13}$$

The constraints 9.12 enforce that each row is covered by at least one column.

ACO can be applied in a very straightforward way to the SCP. The columns are chosen as the solution components and have associated a cost and a pheromone trail. The constraints say that each column can be visited by an ant at most once and that a final solution has to cover all rows. A walk of an ant over the graph representation corresponds to the iterative addition of columns to the partial solution obtained so far. Each ant starts with an empty solution and adds columns until a cover is completed. A pheromone trail τ_i and a heuristic information η_i are associated to each column i . A column to be added is chosen with probability

$$p_i^k(t) = \frac{[\tau_i(t)] \cdot [\eta_i]^\beta}{\sum_{l \in \mathcal{N}^k} [\tau_l(t)] \cdot [\eta_l]^\beta} \quad \text{if } i \in \mathcal{N}^k \tag{9.14}$$

where \mathcal{N}^k is the feasible neighborhood of ant k which consists of all columns which cover at least one still uncovered row. The heuristic information η_i can be chosen in several different ways. For example, a simple static information could be used, taking into account only the column cost: $\eta_i = 1/b_i$. A more sophisticated approach would be to consider the total number of rows d_i covered by a column i and to set $\eta_i = d_i/b_i$. The heuristic information could also be made dependent on the partial solution y_k of an ant k . In this case, it can be defined as $\eta_i = e_i/b_i$, where e_i is the so-called *cover value*, that is, the number of additional rows covered when adding column i to the current partial solution. In other words, the heuristic information measures the unit cost of covering one additional row.

An ant ends the solution construction when all rows are covered. In a post-optimization step, an ant can remove redundant columns—columns that only cover rows which are also covered by a subset of other columns in the final solution—or apply some additional local search to improve solutions. The pheromone update can be carried out in a standard way as described in earlier sections.

When applying ACO to the SCP we have two main differences with the previously presented applications: (i) pheromone trails are associated only to components and, (ii) the length of the ant’s walks (corresponding to the lengths of the sequences) may differ among the ants and, hence, the solution construction only ends when all the ants have terminated their corresponding walks.

There already exist some first applications of ACO to the SCP. In [1], ACO has been used only as a construction algorithm and the approach has only been tested on some

small SCP instances. A more recent article [55] applies Ant System to the SCP and uses techniques to remove redundant columns and local search to improve solutions. Good results are obtained on a large set of benchmark instances taken from ORLIB, but the performance of Ant System could not fully reach that of the best performing algorithms for the SCP.

Example 9.4. *AntNet for network routing applications*

Given a graph representing a telecommunications network, the problem solved by AntNet is to find the minimum cost path between each pair of vertices of the graph. It is important to note that, although finding a minimum cost path on a graph is an easy problem (it can be efficiently solved by algorithms having polynomial worst case complexity [5]), it becomes extremely difficult when the costs on the edges are time-varying stochastic variables. This is the case of routing in packet-switched networks, the target application for AntNet.

Here we briefly describe a simplified version of AntNet (the interested reader should refer to [26], where the AntNet approach to routing is explained and evaluated in detail). As stated earlier, in AntNet each ant searches for a minimum cost path between a given pair of vertices of the network. To this goal, ants are launched from each network vertex towards randomly selected destination vertices. Each ant has a source vertex s and a destination vertex d , and moves from s to d hopping from one vertex to the next until vertex d is reached. When ant k is in vertex i , it chooses the next vertex j to move to according to a probabilistic decision rule which is a function of the ant's memory and of local pheromone and heuristic information (very much like to what happened, for example, in AS).

Unlike AS, where pheromone trails are associated to edges, in AntNet pheromone trails are associated to edge-destination pairs. That is, each directed edge (i, j) has $n - 1$ trail values $\tau_{ijd} \in [0, 1]$ associated, where n is the number of vertices in the graph associated to the routing problem; in general, $\tau_{ijd} \neq \tau_{jid}$. In other words, there is one trail value τ_{ijd} for each possible destination vertex d an ant located in vertex i can have. Each edge has also associated an heuristic value $\eta_{ij} \in [0, 1]$ independent of the final destination. The heuristic values can be set for example to the values $\eta_{ij} = 1 - q_{ij} / \sum_{l \in \mathcal{N}_i} q_{il}$, where q_{ij} is the length (in bits waiting to be sent) of the queue of the link connecting vertex i with its neighbor j : links with a shorter queue have a higher heuristic value.

In AntNet, as well as in most other implementations of ACO algorithms for routing problems [79,90], the daemon component (see Figure 9.5) is not present.

Ants choose their way probabilistically, using as probability a functional composition of the local pheromone trails τ_{ijd} and of the heuristic values η_{ij} . While building the path to their destinations, ants move using the same link queues as data and experience the same delays as data packets. Therefore, the time T_{sd} elapsed while moving from the source vertex s to the destination vertex d can be used as a measure of the quality of the path they built. The overall quality of a path is evaluated by an heuristic function of the trip time T_{sd} and of a local adaptive statistical model maintained in each vertex. In fact, paths need to be evaluated relative to the network status because a trip time T judged of low quality under low congestion conditions could be an excellent one under high traffic load. Once the generic ant k has completed a path, it deposits on the visited vertices an amount of pheromone $\Delta \tau^k(t)$ proportional to the quality of the path it built. To deposit pheromone, after reaching its destination vertex, the ant moves

back to its source vertex along the same path but backward and using high priority queues, to allow a fast propagation of the collected information. The pheromone trail intensity of each edge l_{ij} the ant used while it was moving from s to d is increased as follows: $\tau_{ij d}(t) \leftarrow \tau_{ij d}(t) + \Delta\tau^k(t)$. After the pheromone trail on the visited edges has been updated, the pheromone value of all the outgoing connections of the same vertex i , relative to the destination d , evaporates in such a way that the pheromone values are normalized and can continue to be usable as probabilities: $\tau_{ij d}(t) \leftarrow \tau_{ij d}(t)/(1 + \Delta\tau^k(t))$, $\forall j \in \mathcal{N}_i$, where \mathcal{N}_i is the set of neighbors of vertex i .

AntNet was compared with many state-of-the-art algorithms on a large set of benchmark problems under a variety of traffic conditions. It always compared favorably with competing approaches and it was shown to be very robust with respect to varying traffic conditions and parameter settings. More details on the experimental results can be found in [26].

Applications of the ACO metaheuristic

ACO has recently raised a lot of interest in the scientific community. There are now available numerous successful implementations of the ACO metaheuristic applied to a wide range of different combinatorial optimization problems. These applications comprise two main application fields.

- **\mathcal{NP} -hard problems**, for which the best known algorithms have exponential time worst case complexity. For these problems, very often ACO algorithms are coupled with extra capabilities such as problem-specific local optimizers, which take the ants' solutions to local optima.
- Shortest path problems in which the properties of the problem's graph representation change over time concurrently with the optimization process that has to adapt to the problem's dynamics. In this case, the problem's graph can be physically available (as in network problems), but its properties, like the costs of components or of connections, can change over time. In this case we conjecture that the use of ACO algorithms becomes more and more appropriate as the variation rate of the costs increases and/or the knowledge about the variation process diminishes.

These applications are summarized in Table 9.1. In some of these applications, ACO algorithms have obtained world-class performance, which is the case, for example, for quadratic assignment [63,88], sequential ordering [45,46], vehicle routing [47], scheduling [22,67] or packet-switched network routing [26].

6 DISCUSSION OF APPLICATION PRINCIPLES

Despite being a rather recent metaheuristic, ACO algorithms have already been applied to a large number of different combinatorial optimization problems. Based on this experience, we have identified some basic issues which play an important role in several of these applications. These are discussed in the following.

6.1 Pheromone Trails Definition

An initial, very important choice when applying ACO is the definition of the intended meaning of the pheromone trails. Let us explain this issue with an example. When

Table 9.1. Current applications of ACO algorithms. Applications are listed by class of problems and in chronological order

Problem name	Authors	Algorithm name	Year	Main references
Traveling salesman	Dorigo, Maniezzo & Colomi	AS	1991	[30,37,38]
	Gambardella & Dorigo	Ant-Q	1995	[43]
	Dorigo & Gambardella	ACS & ACS-3-change	1996	[34,35,44]
	Stützle & Hoos	\mathcal{M} MAS	1997	[84,87,88]
	Bullinheimer, Hartl & Strauss	AS_{rank}	1997	[14]
	Cordón, et al.	BWAS	2000	[19]
Quadratic assignment	Maniezzo, Colomi & Dorigo	AS-QAP	1994	[66]
	Gambardella, Taillard & Dorigo	HAS-QAP ^a	1997	[48]
	Stützle & Hoos	\mathcal{M} MAS-QAP	1997	[81,88]
	Maniezzo	ANTS-QAP	1998	[63]
	Maniezzo & Colomi	AS-QAP ^b	1999	[65]
	Colomi, Dorigo & Maniezzo	AS-JSP	1994	[18]
Scheduling problems	Stützle	\mathcal{M} MAS-FSP	1997	[82]
	Bauer et al.	ACS-SMTTP	1999	[2]
	den Besten, Stützle & Dorigo	ACS-SMTWTP	1999	[22]
	Merkle, Middendorf & Schmeck	ACO-RCPS	2000	[67]
	Bullinheimer, Hartl & Strauss	AS-VRP	1997	[12,13]
	Gambardella, Taillard & Agazzi	HAS-VRP	1999	[47]

Connection-oriented network routing	Schoonderwoerd et al.	ABC	1996	[78,79]
	White, Pagurek & Oppacher	ASGA	1998	[91]
	Di Caro & Dorigo	AntNet-FS	1998	[27]
	Bonabeau et al.	ABC-smart ants	1998	[10]
Connection-less	Di Caro & Dorigo	AntNet & AntNet-FA	1997	[24,26,29]
	Heusse et al.	CAF	1998	[56]
	van der Put & Rothkrantz	ABC-backward	1998	[90]
Sequential ordering	Gambardella & Dorigo	HAS-SOP	1997	[45,46]
Graph coloring	Costa & Hertz	ANTCOL	1997	[20]
Shortest common supersequence	Michel & Middendorf	AS-SCS	1998	[69,70]
Frequency assignment	Maniezzo & Carbonaro	ANTS-FAP	1998	[64]
Generalized assignment	Ramalhinho Lourenço & Serra	M/MAS-GAP	1998	[75]
Multiple knapsack	Leguizamón & Michalewicz	AS-MKP	1999	[61]
Optical networks routing	Navarro Varela & Sinclair	ACO-VWP	1999	[73]
Redundancy allocation	Liang & Smith	ACO-RAP	1999	[62]
Constraint satisfaction	Solnon	Ant-P-solver	2000	[80]

^a HAS-QAP is an ant algorithm which does not follow all the aspects of the ACO metaheuristic. ^b This is a variant of the original AS-QAP.

applying ACO to the TSP, the standard interpretation of a pheromone trail τ_{ij} , used in all published ACO applications to the TSP, is that it refers to the desirability of visiting city j directly after a city i . That is, it provides some information on the desirability of the relative positioning of city i and j . Yet, another possibility, not working so well in practice, would be to interpret τ_{ij} as the desirability of visiting city i as the j th city in a tour, that is, the desirability of the absolute positioning. Conversely, when applying ACO to the SMTWTP (see Section 5) better results are obtained when using the absolute position interpretation of the pheromone trails, where τ_{ij} is the desirability of putting job j on the i th position [21]. This is intuitively due to the different role that permutations have in the two problems. In the TSP, permutations are cyclic, that is, only the relative order of the solution components is important and a permutation $\pi = (1\ 2 \dots n)$ has the same tour length as the permutation $\pi' = (n\ 1\ 2 \dots n - 1)$ —it represents the same tour. Therefore, a relative position based pheromone trail is the appropriate choice. On the contrary, in the SMTWTP (as well as in many other scheduling problems), π and π' represent two different solutions with most probably very different costs. Hence, in the SMTWTP the absolute position based pheromone trails are a better choice. Nevertheless, it should be noted that, in principle, both choices are possible, because any solution of the search space can be generated with both representations.

The definition of the pheromone trails is crucial and a poor choice at this stage of the algorithm design will result in poor performance. Fortunately, for many problems the intuitive choice is also a very good one, as it was the case for the previous example applications. Yet, sometimes the use of the pheromones can be somewhat more involved, which is, for example, the case in the ACO application to the shortest common supersequence problem [70].

6.2 Balancing Exploration and Exploitation

Any effective metaheuristic algorithm has to achieve an appropriate balance between the exploitation of the search experience gathered so far and the exploration of unvisited or relatively unexplored search space regions. In ACO several ways exist of achieving such a balance, typically through the management of the pheromone trails. In fact, the pheromone trails induce a probability distribution over the search space and determine which parts of the search space are effectively sampled, that is, in which part of the search space the constructed solutions are located with higher frequency. Note that, depending on the distribution of the pheromone trails, the sampling distribution can vary from a uniform distribution to a degenerate distribution which assigns probability one to a single solution and zero probability to all the others. In fact, this latter situation corresponds to the stagnation of the search as explained on page 261.

The simplest way to exploit the ants' search experience is to make the pheromone update a function of the solution quality achieved by each particular ant. Yet, this bias alone is often too weak to obtain good performance, as was shown experimentally on the TSP [84, 88]. Therefore, in many ACO algorithms (see Section 4) an *elitist strategy* was introduced whereby the best solutions found during the search strongly contribute to pheromone trail updating.

A stronger exploitation of the “learned” pheromone trails can also be achieved during solution construction by applying the pseudo-random proportional rule of Ant Colony System, as explained in Section 4.2.2.

Search space exploration is achieved in ACO primarily by the ants' randomized solution construction. Let us consider for a moment an ACO algorithm that does not use heuristic information (this can be easily achieved by setting $\beta = 0$). In this case, the pheromone updating activity of the ants will cause a shift from the initial uniform sampling of the search space to a sampling focused on specific search space regions. Hence, exploration of the search space will be higher in the initial iterations of the algorithm, and will decrease as the computation goes on. Obviously, attention must be paid to avoid too strong a focus on apparently good regions of the search space, which can cause the ACO algorithm to enter a stagnation situation.

There are several ways to try to avoid such stagnation situations, thus maintaining a reasonable level of exploration of the search space. For example, in ACS the ants use a local pheromone update rule during the solution construction to make the path they have taken less desirable for following ants and, thus, to diversify search, *MMAS* introduces an explicit lower limit on the pheromone trail level so that a minimal level of exploration is always guaranteed. *MMAS* also uses a reinitialization of the pheromone trails, which is a way of enforcing search space exploration. Experience has shown that pheromone trail reinitialization, when combined with appropriate choices for the pheromone trail update [88], can be very useful to refocus the search on a different search space region.

Finally, an important, though somewhat neglected, role in the balance of exploration and exploitation is that of the parameters α and β , which determine the relative influence of pheromone trail and heuristic information. Consider first the influence of parameter α . For $\alpha > 0$, the larger the value of α the stronger the exploitation of the search experience, for $\alpha = 0$ the pheromone trails are not taken into account at all, and for $\alpha < 0$ the most probable choices taken by the ants are those that are less desirable from the point of view of pheromone trails. Hence, varying α could be used to shift from exploration to exploitation and vice versa. The parameter β determines the influence of the heuristic information in a similar way. In fact, systematic variations of α and β could, similarly to what is done in the strategic oscillations approach [50], be part of simple and useful strategies to balance exploration and exploitation.

6.3 ACO and Local Search

In many applications to \mathcal{NP} -hard combinatorial optimization problems, ACO algorithms perform best when coupled with local search algorithms (which is, in fact, a particular type of daemon action of the ACO metaheuristic). Local search algorithms locally optimize the ants' solutions and these locally optimized solutions are used in the pheromone update.

The use of local search in ACO algorithms can be very interesting as the two approaches are complementary. In fact, ACO algorithms perform a rather coarse-grained search, and the solutions they produce can then be locally optimized by an adequate local search algorithm. The coupling can therefore greatly improve the quality of the solutions generated by the ants.

On the other side, generating initial solutions for local search algorithms is not an easy task. For example, it has been shown that, for most problems, repeating local searches from randomly generated initial solutions is not efficient (see, e.g., [57]). In practice, ants probabilistically combine solution components which are part of the best locally optimal solutions found so far and generate new, promising initial solutions

for the local search. Experimentally, it has been found that such a combination of a probabilistic, adaptive construction heuristic with local search can yield excellent results [6,35,87].

It is important to note that when using local search a choice must be made concerning pheromone trail update: either pheromone is added to the components and/or connections of the locally optimal solution, or to the starting solutions for the local search.¹¹ The quasi totality of published research has used the first approach. Although it could be interesting to investigate the second approach, some recent experimental results suggest that its performance is worse.

Despite the fact that the use of local search algorithms has been shown to be crucial for achieving best performance in many ACO applications, it should be noted that ACO algorithms also show very good performance where local search algorithms cannot be applied easily. One such example are the network routing applications described in Section 5 or the shortest common supersequence problem [70].

6.4 Heuristic Information

The possibility of using heuristic information to direct the ants' probabilistic solution construction is important because it gives the possibility of exploiting problem specific knowledge. This knowledge can be available *a priori* (this is the most frequent situation in static problems) or at run-time (this is the typical situation in dynamic problems). In static problems, the heuristic information η is usually computed once at initialization time and then it remains the same throughout the whole algorithm's run. An example is the use, in the TSP applications, of the length d_{ij} of the edge connecting cities i and j to define the heuristic information $\eta_{ij} = 1/d_{ij}$. Static heuristic information has the advantage that (i) it is easy to compute, (ii) it has to be computed only once at initialization time, and (iii) in each iteration of the ACO algorithm, a table can be pre-computed with the values of $\tau_{ij}(t) \cdot \eta_{ij}^\beta$, which can result in a very significant saving of computation time. In the dynamic case, the heuristic information may also depend on the partial solution constructed so far and therefore has to be computed at each step of an ant's walk. This determines a higher computational cost that may be compensated by the higher accuracy of the computed heuristic values. For example, in the ACO application to the SMTWTP we found that the use of dynamic heuristic information based on the modified due date or the apparent urgency heuristics (see Section 5) resulted in a better overall performance.

Another way of computing heuristic information was introduced in the ANTS algorithm [63], where it is computed using lower bounds on the solution cost of the completion of an ant's partial solution. This method has the advantage that it facilitates the exclusion of certain choices because they lead to solutions that are worse than the best found so far. It allows therefore the combination of knowledge on the calculation of lower bounds from mathematical programming with the ACO paradigm. Nevertheless, a disadvantage is that the computation of the lower bounds can be time consuming, especially because they have to be calculated at each single construction step by each ant.

¹¹ Making an analogy with genetic algorithms, we could call the first approach "Lamarckian" and the second "Darwinian".

Finally, it should be noted that while the use of heuristic information is rather important for a generic ACO algorithm, its importance is strongly reduced if local search is used to improve solutions. This is due to the fact that local search takes into account information about the cost to improve solutions in a more direct way. Fortunately, this means that ACO algorithms can also achieve, in combination with a local search algorithm, very good performance for problems for which it is difficult to define *a priori* a very informative heuristic information.

6.5 Number of Ants

Why use a colony of ants instead of using one single ant? In fact, although a single ant is capable of generating a solution, efficiency considerations suggest that the use of a colony of ants is often a desirable choice. This is particularly true for geographically distributed problems, because the differential path length effect exploited by ants in the solution of this class of problems can only arise in presence of a colony of ants. It is also interesting to note that in routing problems ants solve many shortest path problems in parallel (one between each pair of vertices) and a colony of ants must be used for each of these problems.

On the other hand, in the case of combinatorial optimization problems the differential length effect is not exploited and the use of m ants, $m > 1$, that build r solutions each (i.e., the ACO algorithm is run for r iterations) could be equivalent to the use of one ant that generates $m \cdot r$ solutions. Nevertheless, in this case theoretical results on the convergence of some specific ACO algorithms, which will be presented in Section 7, as well as experimental evidence suggest that ACO algorithms perform better when the number m of ants is set to a value $m > 1$.

In general, the best value for m is a function of the particular ACO algorithm chosen as well as of the class of problems being attacked, and most of the times it must be set experimentally. Fortunately, ACO algorithms seem to be rather robust to the actual number of ants used.

6.6 Candidate Lists

One possible difficulty encountered by ACO algorithms is when they are applied to problems with a large neighborhood in the solution construction. In fact, an ant that visits a state with a large neighborhood has a correspondingly large number of possible moves among which to choose. Possible problems are that the solution construction is significantly slowed down and that the probability that many ants visit the same state is very small. Such a situation can occur, for example, in the ACO application to large TSPs or large SCPs.

In such situations, the above-mentioned problem can be considerably reduced by the use of candidate lists. Candidate lists comprise a small set of promising neighbors of the current state. They are created using *a priori* available knowledge on the problem, if available, or dynamically generated information. Their use allows ACO algorithms to focus on the more interesting components, strongly reducing the dimension of the search space.

As an example, consider the ACO application to the TSP. For the TSP it is known that very often optimal solutions can be found within a surprisingly small subgraph consisting of all the cities and of those edges that connect each city to only a few of its nearest neighbors. For example, for the TSPLIB instance `pr2392.tsp` with 2392 cities

an optimal solution can be found within a subgraph of the 8 nearest neighbors [76]. This knowledge can be used for defining candidate lists, which was first done in the context of ACO algorithms in [44]. A candidate list includes for each city its cl nearest neighbors. During solution construction an ant tries to choose the city to move to only among the cities in the candidate list. Only if all these cities have already been visited, the ant can choose among the other cities.

So far, in ACO algorithms the use of candidate lists or similar approaches is still rather unexplored. Inspiration from other techniques like Tabu Search [51] or GRASP [42], where strong use of candidate lists is made, could be useful for the development of effective candidate list strategies for ACO.

7 OTHER DEVELOPMENTS

7.1 Proving Convergence

The simplest stochastic optimization algorithm is random search. Besides simplicity, random search has also the nice property that it guarantees that it will find, sooner or later, the optimal solution to your problem. Unfortunately, it is very inefficient. Stochastic optimization algorithms can be seen as ways of biasing random search so to make it more efficient. Unfortunately, once a stochastic algorithm is biased, it is no longer guaranteed to find, at some point, the optimal solution. In fact, the bias could simply rule out this possibility. It is therefore interesting to have convergence proofs that reassure you that this does not happen.

Although the problem of proving convergence to the optimal solution of a generic ACO algorithm is open (and it will most probably remain so, given the generality of the ACO metaheuristic), convergence has recently been proved for a few instances of ACO algorithms. The first of such proofs was provided by Gutjahr [53,54] who proved convergence to the optimal solution for a particular ACO algorithm he calls *Graph-based Ant System* (GBAS): the proof states that, given a small $\epsilon > 0$ and for fixed values of some algorithm parameters, after a number of cycles $t \geq t_0$ the algorithm will find the optimal solution with probability $P_t \geq 1 - \epsilon$, where $t_0 = f(\epsilon)$. A limitation of this important result is that it applies to an ACO algorithm, GBAS, for which no experimental results are available (i.e., we do not know what is its performance). More recently, Stützle and Dorigo [86] have proved convergence for two of the experimentally most successful ACO algorithms: ACS and *MMAS*.

7.2 Parallel Implementations

The very nature of ACO algorithms lends them to be parallelized in the data or population domains. In particular, many parallel models used in other population-based algorithms can be easily adapted to the ACO structure. Most parallelization strategies can be classified into *fine-grained* and *coarse-grained* strategies. Characteristic of fine-grained parallelization is that very few individuals are assigned to one single processor and that frequent information exchange among the processors takes place. On the contrary, in coarse grained approaches larger subpopulations or even full populations are assigned to single processors and information exchange is rather rare. We refer, for example, to [16] for an overview.

Fine-grained parallelization schemes have been investigated with parallel versions of AS for the TSP on the Connection Machine CM-2 adopting the approach of attributing a single processing unit to each ant [7]. Experimental results showed that communication overhead can be a major problem with this approach on fine grained parallel machines, since ants end up spending most of their time communicating to other ants the modifications they made to pheromone trails. Similar negative results have also been reported in [15].

As shown by several researches [7,15,59,71,83], coarse grained parallelization schemes are much more promising for ACO. When applied to ACO, coarse grained schemes run p subcolonies in parallel, where p is the number of available processors. Information among the subcolonies is exchanged at certain intervals. For example, in the Partially Asynchronous Parallel Implementation (PAPI) of Bullheimer, Kotsis and Strauss [15], for which high speed-up was observed, the subcolonies exchange pheromone information every fixed number of iterations performed by each subcolony. Krüger, Merkle and Middendorf [59] investigated which information should be exchanged between the m subcolonies and how this information should be used to update the subcolony's trail information. Their results showed that it was better to exchange the best solutions found so far and to use them in the pheromone update than to exchange complete pheromone matrices for modifications of the pheromone matrix of a local subcolony. Middendorf, Reischle, and Schmeck [71] investigate different ways of exchanging solutions among m ant colonies. They consider an exchange of the global best solutions among all colonies and local exchanges based on a virtual neighborhood among subcolonies which corresponds to a directed ring. Their main observation was that the best solutions, with respect to computing time and solution quality, were obtained by limiting the information exchange to a local neighborhood of the colonies. In the extreme case, there is no communication among the subcolonies, resulting in parallel independent runs of an algorithm. This is the easiest way to parallelize randomized algorithms and can be very effective as has been shown by computational results presented by Stützle [83].

8 CONCLUSIONS

The field of ACO algorithms is very lively, as testified, for example, by the successful biannual workshop (ANTS—From Ant Colonies to Artificial Ants: A Series of International Workshops on Ant Algorithms; <http://iridia.ulb.ac.be/~ants/>) where researchers meet to discuss the properties of ACO and other ant algorithms, both theoretically and experimentally.

From the theory side, the most interesting ongoing work concerns the study of the relationship between ACO algorithms and other well-known stochastic optimization techniques. For example, it has been shown [68] that, when interpreting ACO algorithms as methods for searching in the space of pheromones (i.e., artificial ants are seen as procedures to update pheromone trails so to increase the probability of generating very good solutions to the combinatorial optimization problem), then AS can be interpreted as an approximate form of stochastic gradient descent [92] (a well-known algorithm which has been extensively used in machine learning) in the space of pheromones. Similarly, it has been shown [41,94] that there are strong connections between ACO algorithms and the Cross-Entropy method [77].

From the experimental side, most of the current research is in the direction of increasing the number of problems that are successfully solved by ACO algorithms, including real-world, industrial applications [39].

Currently, the great majority of problems attacked by ACO are static and well-defined combinatorial optimization problems, that is, problems for which all the necessary information is available and does not change during problem solution. For this kind of problems ACO algorithms must compete with very well established algorithms, often specialized for the given problem. Also, very often the role played by local search is extremely important to obtain good results (see for example [46]). Although rather successful on these problems, we believe that ACO algorithms will really show their greatest advantage when they are systematically applied to “ill-structured” problems for which it is not clear how to apply local search, or to highly dynamic domains with only local information available. A first step in this direction has already been made with the application to telecommunications networks routing, but much further research will be necessary. The reader interested in learning more about ACO is referred to the book “Ant Colony Optimization” by the same authors [40].

ACKNOWLEDGMENTS

We thank Joshua Knowles and Michael Sampels for their useful comments on a draft version of this paper. Marco Dorigo acknowledges support from the Belgian FNRS, of which he is a Senior Research Associate. This work was partially supported by the “Metaheuristics Network”, a Research Training Network funded by the Improving Human Potential programme of the CEC, grant HPRN-CT-1999-00106. The information provided is the sole responsibility of the authors and does not reflect the Community’s opinion. The Community is not responsible for any use that might be made of data appearing in this publication.

REFERENCES

- [1] D.A. Alexandrov and Y.A. Kochetov (2000) The behavior of the ant colony algorithm for the set covering problem. In: K. Inderfurth, G. Schwödiauer, W. Domschke, F. Juhnke, P. Kleinschmidt, and G. Wäscher (eds.), *Operations Research Proceedings 1999*. Springer-Verlag, Berlin, Germany, pp. 255–260.
- [2] A. Bauer, B. Bullnheimer, R.F. Hartl and C. Strauss (1999) An ant colony optimization approach for the single machine total tardiness problem. In: *Proceedings of the 1999 Congress on Evolutionary Computation (CEC’99)*, IEEE Press, Piscataway, NJ, pp. 1445–1450.
- [3] R. Beckers, J.-L. Deneubourg and S. Goss (1993) Modulation of trail laying in the ant *Lasius niger* (hymenoptera: Formicidae) and its role in the collective selection of a food source. *Journal of Insect Behavior*, **6**(6), 751–759.
- [4] R. Bellman, A.O. Esogbue and I. Nabeshima (1982) *Mathematical Aspects of Scheduling and Applications*. Pergamon Press, New York, NJ.
- [5] D. Bertsekas (1998) *Network Optimization: Continuous and Discrete Models*. Athena Scientific, Belmont, MA.

- [6] K.D. Boese, A.B. Kahng and S. Muddu (1994) A new adaptive multi-start technique for combinatorial global optimization. *Operations Research Letters*, **16**, 101–113.
- [7] M. Bolondi and M. Bondanza (1993) Parallelizzazione di un algoritmo per la risoluzione del problema del commesso viaggiatore. Master's thesis, Dipartimento di Elettronica, Politecnico di Milano, Italy.
- [8] E. Bonabeau, M. Dorigo and G. Theraulaz (1999) *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, New York, NJ.
- [9] E. Bonabeau, M. Dorigo and G. Theraulaz (2000) Inspiration for optimization from social insect behavior. *Nature*, **406**, 39–42.
- [10] E. Bonabeau, F. Henaux, S. Guérin, D. Snyers, P. Kuntz and G. Theraulaz (1998) Routing in telecommunication networks with “Smart” ant-like agents. In: *Proceedings of IATA'98, Second International Workshop on Intelligent Agents for Telecommunication Applications*, volume 1437 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, Berlin, Germany, pp. 60–72.
- [11] E. Bonabeau and G. Theraulaz (2000) Swarm smarts. *Scientific American*, **282**(3), 54–61.
- [12] B. Bullnheimer, R.F. Hartl and C. Strauss (1999) Applying the Ant System to the vehicle routing problem. In: S. Voß, S. Martello, I.H. Osman and C. Roucairol (eds.), *Meta-Heuristics: Advances and Trends in Local Search Paradigms for Optimization*. Kluwer Academic Publishers, Dordrecht, The Netherlands, pp. 285–296.
- [13] B. Bullnheimer, R.F. Hartl and C. Strauss (1999) An improved Ant System algorithm for the vehicle routing problem. *Annals of Operations Research*, **89**, 319–328.
- [14] B. Bullnheimer, R.F. Hartl and C. Strauss (1999) A new rank-based version of the Ant System: A computational study. *Central European Journal for Operations Research and Economics*, **7**(1), 25–38.
- [15] B. Bullnheimer, G. Kotsis and C. Strauss (1998) Parallelization strategies for the Ant System. In: R. De Leone, A. Murli, P. Pardalos and G. Toraldo (eds.), *High Performance Algorithms and Software in Nonlinear Optimization*, volume 24 of *Applied Optimization*. Kluwer Academic Publishers, Dordrecht, The Netherlands, pp. 87–100.
- [16] E. Cantú-Paz (2000) *Efficient and Accurate Parallel Genetic Algorithms*. Kluwer Academic Publishers, Boston, MA.
- [17] A. Colomi, M. Dorigo and V. Maniezzo (1992) Distributed optimization by ant colonies. In: F.J. Varela and P. Bourguine (eds.), *Proceedings of the First European Conference on Artificial Life*, MIT Press, Cambridge, MA, pp. 134–142.
- [18] A. Colomi, M. Dorigo, V. Maniezzo and M. Trubian (1994) Ant System for job-shop scheduling. *JORBEL—Belgian Journal of Operations Research, Statistics and Computer Science*, **34**(1), 39–53.
- [19] O. Cordon, I. Fernández de Viana, F. Herrera and L. Moreno (2000) A new ACO model integrating evolutionary computation concepts: The best-worst Ant System. In: M. Dorigo, M. Middendorf and T. Stützle (eds.), *Abstract proceedings*

- of ANTS2000—*From Ant Colonies to Artificial Ants: A Series of International Workshops on Ant Algorithms*. IRIDIA, Université Libre de Bruxelles, Belgium, pp. 22–29.
- [20] D. Costa and A. Hertz (1997) Ants can colour graphs. *Journal of the Operational Research Society*, **48**, 295–305.
- [21] M. den Besten (2000) Ants for the single machine total weighted tardiness problem. Master's thesis, University of Amsterdam, The Netherlands.
- [22] M.L. den Besten, T. Stützle and M. Dorigo (2000) Ant colony optimization for the total weighted tardiness problem. In: M. Schoenauer, K. Deb, G. Rudolph, X. Yao, E. Lutton, J.J. Merelo and H.-P. Schwefel (eds.), *Proceedings of PPSN-VI, Sixth International Conference on Parallel Problem Solving from Nature*, volume 1917 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, Germany, pp. 611–620.
- [23] J.-L. Deneubourg, S. Aron, S. Goss and J.-M. Pasteels (1990) The self-organizing exploratory pattern of the Argentine ant. *Journal of Insect Behavior*, **3**, 159–168.
- [24] G. Di Caro and M. Dorigo (1997) AntNet: A mobile agents approach to adaptive routing. Technical Report IRIDIA/97-12, IRIDIA, Université Libre de Bruxelles, Belgium.
- [25] G. Di Caro and M. Dorigo (1998) Ant colonies for adaptive routing in packet-switched communications networks. In: A.E. Eiben, T. Bäck, M. Schoenauer and H.-P. Schwefel (eds.), *Proceedings of PPSN-V, Fifth International Conference on Parallel Problem Solving from Nature*, volume 1498 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, Germany, pp. 673–682.
- [26] G. Di Caro and M. Dorigo (1998) AntNet: Distributed stigmergetic control for communications networks. *Journal of Artificial Intelligence Research*, **9**, 317–365.
- [27] G. Di Caro and M. Dorigo (1998) Extending AntNet for best-effort Quality-of-Service routing. Unpublished presentation at ANTS'98—*From Ant Colonies to Artificial Ants: First International Workshop on Ant Colony Optimization*, <http://iridia.ulb.ac.be/ants98/ants98.html>, October 15–16.
- [28] G. Di Caro and M. Dorigo (1998) Mobile agents for adaptive routing. In: H. El-Rewini (ed.), *Proceedings of the 31st International Conference on System Sciences (HICSS-31)*. IEEE Computer Society Press, Los Alamitos, CA, pp. 74–83.
- [29] G. Di Caro and M. Dorigo (1998) Two ant colony algorithms for best-effort routing in datagram networks. In: Y. Pan, S.G. Akl and K. Li (eds.), *Proceedings of the Tenth IASTED International Conference on Parallel and Distributed Computing and Systems (PDCS'98)*, IASTED/ACTA Press, Anaheim, CA, pp. 541–546.
- [30] M. Dorigo (1992) *Optimization, Learning and Natural Algorithms* (in Italian). PhD thesis, Dipartimento di Elettronica, Politecnico di Milano, Italy, 140 pp.
- [31] M. Dorigo, E. Bonabeau and G. Theraulaz (2000) Ant algorithms and stigmergy. *Future Generation Computer Systems*, **16**(8), 851–871.

- [32] M. Dorigo and G. Di Caro (1999) The Ant Colony Optimization meta-heuristic. In: D. Corne, M. Dorigo and F. Glover (eds.), *New Ideas in Optimization*. McGraw Hill, London, UK, pp. 11–32.
- [33] M. Dorigo, G. Di Caro and L.M. Gambardella (1999) Ant algorithms for discrete optimization. *Artificial Life*, **5**(2), 137–172.
- [34] M. Dorigo and L.M. Gambardella (1997) Ant colonies for the traveling salesman problem. *BioSystems*, **43**, 73–81.
- [35] M. Dorigo and L.M. Gambardella (1997) Ant Colony System: A cooperative learning approach to the traveling salesman problem. *IEEE Transactions on Evolutionary Computation*, **1**(1), 53–66.
- [36] M. Dorigo, V. Maniezzo and A. Colorni (1991) The Ant System: An autocatalytic optimizing process. Technical Report 91-016 Revised, Dipartimento di Elettronica, Politecnico di Milano, Italy.
- [37] M. Dorigo, V. Maniezzo and A. Colorni (1991) Positive feedback as a search strategy. Technical Report 91-016, Dipartimento di Elettronica, Politecnico di Milano, Italy.
- [38] M. Dorigo, V. Maniezzo, and A. Colorni (1996) The Ant System: Optimization by a colony of cooperating agents. *IEEE Transactions on Systems, Man, and Cybernetics—Part B*, **26**(1), 29–41.
- [39] M. Dorigo, M. Middendorf and T. Stützle (2000) (eds.) *Abstract proceedings of ANTS2000—From Ant Colonies to Artificial Ants: A Series of International Workshops on Ant Algorithms*. IRIDIA, Université Libre de Bruxelles, Belgium, 7–9 September.
- [40] M. Dorigo and T. Stützle *Ant Colony Optimization*. MIT Press, Cambridge, MA (forthcoming).
- [41] M. Dorigo, M. Zlochin, N. Meuleau and M. Birattari (2001) Updating ACO pheromones using stochastic gradient ascent and cross-entropy methods. Technical Report IRIDIA/2001-19, IRIDIA, Université Libre de Bruxelles, Belgium. *Proceedings of the 2nd European Workshop on Evolutionary Computation in Combinatorial Optimization (EvoCOP-2002)*, (to appear).
- [42] T.A. Feo and M.G.C. Resende (1995) Greedy randomized adaptive search procedures. *Journal of Global Optimization*, **6**, 109–133.
- [43] L.M. Gambardella and M. Dorigo (1995) Ant-Q: A reinforcement learning approach to the traveling salesman problem. In: A. Prieditis and S. Russell (eds.), *Proceedings of the Twelfth International Conference on Machine Learning (ML-95)*, Morgan Kaufmann Publishers, Palo Alto, CA, pp. 252–260.
- [44] L.M. Gambardella and M. Dorigo (1996) Solving symmetric and asymmetric TSPs by ant colonies. In: *Proceedings of the 1996 IEEE International Conference on Evolutionary Computation (ICEC'96)*, IEEE Press, Piscataway, NJ, pp. 622–627.
- [45] L.M. Gambardella and M. Dorigo (1997) HAS-SOP: An hybrid Ant System for the sequential ordering problem. Technical Report IDSIA-11-97, IDSIA, Lugano, Switzerland.

- [46] L.M. Gambardella and M. Dorigo (2000) Ant Colony System hybridized with a new local search for the sequential ordering problem. *INFORMS Journal on Computing*, **12**(3), 237–255.
- [47] L.M. Gambardella, È D. Taillard and G. Agazzi (1999) MACS-VRPTW: A multiple ant colony system for vehicle routing problems with time windows. In: D. Corne, M. Dorigo and F. Glover (eds.), *New Ideas in Optimization*. McGraw Hill, London, UK, pp. 63–76.
- [48] L.M. Gambardella, È.D. Taillard and M. Dorigo (1999) Ant colonies for the quadratic assignment problem. *Journal of the Operational Research Society*, **50**(2), 167–176.
- [49] M.R. Garey and D.S. Johnson (1979) *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman, San Francisco, CA.
- [50] F. Glover (1990) Tabu search—part II. *ORSA Journal on Computing*, **2**(1), 4–32.
- [51] F. Glover and M. Laguna (1997) *Tabu Search*. Kluwer Academic Publishers, Boston, MA.
- [52] S. Goss, S. Aron, J.L. Deneubourg and J.M. Pasteels (1989) Self-organized shortcuts in the Argentine ant. *Naturwissenschaften*, **76**, 579–581.
- [53] W.J. Gutjahr (2000) A Graph-based Ant System and its convergence. *Future Generation Computer Systems*, **16**(8), 873–888.
- [54] W.J. Gutjahr (2002) ACO algorithms with guaranteed convergence to the optimal solution. *Information Processing Letters*, (in press).
- [55] R. Hadji, M. Rahoual, E. Talbi and V. Bachelet (2000) Ant colonies for the set covering problem. In: M. Dorigo, M. Middendorf and T. Stützle (eds.), *Abstract proceedings of ANTS2000—From Ant Colonies to Artificial Ants: A Series of International Workshops on Ant Algorithms*. IRIDIA, Université Libre de Bruxelles, Belgium, pp. 63–66.
- [56] M. Heusse, S. Guérin, D. Snyers and P. Kuntz (1998) Adaptive agent-driven routing and load balancing in communication networks. *Advances in Complex Systems*, **1**(2), 237–254.
- [57] D.S. Johnson and L.A. McGeoch (1997) The travelling salesman problem: A case study in local optimization. In: E.H.L. Aarts and J.K. Lenstra (eds.), *Local Search in Combinatorial Optimization*. John Wiley & Sons, Chichester, UK, pp. 215–310.
- [58] M. Jünger, G. Reinelt and S. Thienel (1994) Provably good solutions for the traveling salesman problem. *Zeitschrift für Operations Research*, **40**, 183–217.
- [59] F. Krüger, D. Merkle and M. Middendorf Studies on a parallel Ant System for the BSP model. Unpublished manuscript.
- [60] E.L. Lawler, J.K. Lenstra, A.H.G. Rinnooy Kan and D.B. Shmoys (1985) *The Travelling Salesman Problem*. John Wiley & Sons, Chichester, UK.
- [61] G. Leguizamón and Z. Michalewicz (1999) A new version of Ant System for subset problems. In: *Proceedings of the 1999 Congress on Evolutionary Computation (CEC'99)*. IEEE Press, Piscataway, NJ, pp. 1459–1464.

- [62] Y.-C. Liang and A.E. Smith (1999) An Ant System approach to redundancy allocation. In: *Proceedings of the 1999 Congress on Evolutionary Computation (CEC'99)*. IEEE Press, Piscataway, NJ, pp. 1478–1484.
- [63] V. Maniezzo (1999) Exact and approximate nondeterministic tree-search procedures for the quadratic assignment problem. *INFORMS Journal on Computing*, **11**(4), 358–369.
- [64] V. Maniezzo and A. Carbonaro (2000) An ANTS heuristic for the frequency assignment problem. *Future Generation Computer Systems*, **16**(8), 927–935.
- [65] V. Maniezzo and A. Colomi (1999) The Ant System applied to the quadratic assignment problem. *IEEE Transactions on Data and Knowledge Engineering*, **11**(5), 769–778.
- [66] V. Maniezzo, A. Colomi and M. Dorigo (1994) The Ant System applied to the quadratic assignment problem. Technical Report IRIDIA/94-28, IRIDIA, Université Libre de Bruxelles, Belgium.
- [67] D. Merkle, M. Middendorf and H. Schmeck (2000) Ant colony optimization for resource-constrained project scheduling. In: *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2000)*, Morgan Kaufmann Publishers, San Francisco, CA, pp. 893–900.
- [68] N. Meuleau and M. Dorigo (2002) Ant colony optimization and stochastic gradient descent. *Artificial Life*, (in press).
- [69] R. Michel and M. Middendorf (1998) An island model based Ant System with lookahead for the shortest supersequence problem. In: A.E. Eiben, T. Bäck, M. Schoenauer and H.-P. Schwefel (eds.), *Proceedings of PPSN-V, Fifth International Conference on Parallel Problem Solving from Nature*, volume 1498 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, Germany, pp. 692–701.
- [70] R. Michel and M. Middendorf (1999) An ACO algorithm for the shortest supersequence problem. In: D. Corne, M. Dorigo and F. Glover (eds.), *New Ideas in Optimization*. McGraw Hill, London, UK, pp. 51–61.
- [71] M. Middendorf, F. Reischle and H. Schmeck (2002) Multi colony ant algorithms. *Journal of Heuristics*, (in press).
- [72] T.E. Morton, R.M. Rachamadugu and A. Vepsalainen (1984) Accurate myopic heuristics for tardiness scheduling. GSIA Working Paper 36-83-84, Carnegie–Mellon University, PA.
- [73] G. Navarro Varela and M.C. Sinclair (1999) Ant colony optimisation for virtual-wavelength-path routing and wavelength allocation. In: *Proceedings of the 1999 Congress on Evolutionary Computation (CEC'99)*. IEEE Press, Piscataway, NJ, pp. 1809–1816.
- [74] C.H. Papadimitriou (1994) *Computational Complexity*. Addison-Wesley, Reading, MA.
- [75] H. Ramalhinho Lourenço and D. Serra (1998) Adaptive approach heuristics for the generalized assignment problem. Technical Report Economic Working Papers Series No. 304, Universitat Pompeu Fabra, Department of Economics and Management, Barcelona, Spain.

- [76] G. Reinelt (1994) *The Traveling Salesman: Computational Solutions for TSP Applications*, volume 840 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, Germany.
- [77] R.Y. Rubinstein (2001) Combinatorial optimization via the simulated cross-entropy method. In: *Encyclopedia of Operations Research and Management Science*. Kluwer Academic Publishers, Boston, MA.
- [78] R. Schoonderwoerd, O. Holland and J. Bruten (1997) Ant-like agents for load balancing in telecommunications networks. In: *Proceedings of the First International Conference on Autonomous Agents*. ACM Press, New York, NY, pp. 209–216.
- [79] R. Schoonderwoerd, O. Holland, J. Bruten and L. Rothkrantz (1996) Ant-based load balancing in telecommunications networks. *Adaptive Behavior*, **5**(2), 169–207.
- [80] C. Solnon (2000) Solving permutation constraint satisfaction problems with artificial ants. In: W. Horn (ed.), *Proceedings of the 14th European Conference on Artificial Intelligence*. IOS Press, Amsterdam, The Netherlands, pp. 118–122.
- [81] T. Stützle (1997) *MA_X-MIN* Ant System for the quadratic assignment problem. Technical Report AIDA-97-4, FG Intellektik, FB Informatik, TU Darmstadt, Germany.
- [82] T. Stützle (1998) An ant approach to the flow shop problem. In: *Proceedings of the 6th European Congress on Intelligent Techniques & Soft Computing (EUFIT'98)*, volume 3, Verlag Mainz, Wissenschaftsverlag, Aachen, Germany, pp. 1560–1564.
- [83] T. Stützle (1998) Parallelization strategies for ant colony optimization. In: A.E. Eiben, T. Bäck, M. Schoenauer and H.-P. Schwefel (eds.), *Proceedings of PPSN-V, Fifth International Conference on Parallel Problem Solving from Nature*, volume 1498 of *Lecture Notes in Computer Science*. Springer Verlag, Berlin, Germany, pp. 722–731.
- [84] T. Stützle (1999) *Local Search Algorithms for Combinatorial Problems: Analysis, Improvements and New Applications*. Infix, Sankt Augustin, Germany.
- [85] T. Stützle and M. Dorigo (1999) ACO algorithms for the quadratic assignment problem. In: D. Corne, M. Dorigo and F. Glover (eds.), *New Ideas in Optimization*. McGraw Hill, London, UK, pp. 33–50.
- [86] T. Stützle and M. Dorigo (2002) A short convergence proof for a class of ACO algorithms. *IEEE Transactions on Evolutionary Computation* (in press).
- [87] T. Stützle and H.H. Hoos (1997) The *MA_X-MIN* Ant System and local search for the traveling salesman problem. In: T. Bäck, Z. Michalewicz and X. Yao (eds.), *Proceedings of the 1997 IEEE International Conference on Evolutionary Computation (ICEC'97)*, IEEE Press, Piscataway, NJ, pp. 309–314.
- [88] T. Stützle and H.H. Hoos (2000) *MA_X-MIN* Ant System. *Future Generation Computer Systems*, **16**(8):889–914.
- [89] R.S. Sutton and A.G. Barto (1998) *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA.
- [90] R. van der Put (1998) Routing in the faxfactory using mobile agents. Technical Report R&D-SV-98-276, KPN Research, The Netherlands.

- [91] T. White, B. Pagurek and F. Oppacher (1998) Connection management using adaptive mobile agents. In: H.R. Arabnia (ed.), *Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA'98)*. CSREA Press, pp. 802–809.
- [92] R.J. Williams (1992) Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, **8**(3), 229–256.
- [93] M. Yannakakis (1997) Computational complexity. In: E.H.L. Aarts and J.K. Lenstra (eds.), *Local Search in Combinatorial Optimization*. John Wiley & Sons, Chichester, UK, pp. 19–55.
- [94] M. Zlochin, M. Birattari, N. Meuleau and M. Dorigo (2001) Combinatorial optimization using model-based search. Technical Report IRIDIA/2001-15, IRIDIA, Université Libre de Bruxelles, Belgium, (submitted to *Annals of Operations Research*).