

A Short Convergence Proof for a Class of Ant Colony Optimization Algorithms

Thomas Stützle and Marco Dorigo, *Senior Member, IEEE*

Abstract—In this paper, we prove some convergence properties for a class of ant colony optimization algorithms. In particular, we prove that for any small constant $\epsilon > 0$ and for a sufficiently large number of algorithm iterations t , the probability of finding an optimal solution at least once is $P^*(t) \geq 1 - \epsilon$ and that this probability tends to 1 for $t \rightarrow \infty$. We also prove that, after an optimal solution has been found, it takes a finite number of iterations for the pheromone trails associated to the found optimal solution to grow higher than any other pheromone trail and that, for $t \rightarrow \infty$, any fixed ant will produce the optimal solution during the t th iteration with probability $P \geq 1 - \hat{\epsilon}(\tau_{\min}, \tau_{\max})$, where τ_{\min} and τ_{\max} are the minimum and maximum values that can be taken by pheromone trails.

Index Terms—ACO algorithms, ant algorithms, ant colony optimization, approximation algorithms, convergence proof, heuristics, metaheuristics.

I. INTRODUCTION

ANT COLONY optimization (ACO) is a metaheuristic for the approximate solution of combinatorial optimization problems that has been inspired by the foraging behavior of ant colonies. In ACO algorithms, the computational resources are allocated to a set of relatively simple agents (artificial ants) that exploit stigmergic communication, i.e., a form of indirect communication mediated by the environment [4], [11] to construct solutions to the considered problem. The construction of good solutions is a result of the agents' cooperative interaction.

In the last ten years, a number of applications to many different \mathcal{NP} -hard combinatorial optimization problems [5], [6] has empirically shown the effectiveness of ant colony optimization. Still, very little theory is available to explain the reasons underlying ACO's success. Birattari *et al.* [1] have proposed an interpretation of ACO in the framework of optimal control and reinforcement learning, while Meuleau and Dorigo [16] have shown that ACO algorithms and stochastic gradient descent are strongly related and that a particular form of ACO algorithms converges with probability 1 to a local optimum. Closer to the

Manuscript received February 15, 2001; revised September 11, 2001. This work was supported by the "Metaheuristics Network," a Research Training Network funded by the Improving Human Potential Programme of the European Commission under Grant HPRN-CT-1999-00106. The work of M. Dorigo was supported by the Belgian Fonds National de la Recherche Scientifique. 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.

T. Stützle is with the Intellectics Group, Computer Science Department, Darmstadt University of Technology, D-64283 Darmstadt, Germany (email: stuetzle@informatik.tu-darmstadt.de).

M. Dorigo is with Institut de Recherches Interdisciplinaires et de Développements en Intelligence Artificielle, Université Libre de Bruxelles, B-1050 Brussels, Belgium (email: mdorigo@ulb.ac.be).

Publisher Item Identifier 10.1109/TEVC.2002.802444.

work presented in this paper is Gutjahr's convergence proof [14]. He proved convergence to the globally optimal solution with probability $1 - \epsilon$ of a particular ant colony optimization algorithm called graph-based ant system (GBAS). Yet, GBAS is quite different from any implemented ACO algorithm and its empirical performance is unknown.

Differently, in this paper, we present a simple convergence proof that applies directly to at least two of the most (experimentally) successful ACO algorithms: 1) ant colony system (ACS) [7] and 2) $\mathcal{MAX-MIN}$ ant system (\mathcal{MMAS}) [19].

The paper is organized as follows. In the next section, we give a formal characterization of the class of problems and of the algorithm for which the convergence proof, presented in Section III, holds. In Section IV, we discuss the meaning and implications of the proof, with particular attention to its relations with Gutjahr's convergence proof. In Section V, we show that the proof holds for a wider class of ACO algorithms and, in particular, for \mathcal{MMAS} and ACS. Additionally, we show that the proposed proof continues to hold if our algorithm is made more general by adding problem-dependent heuristic information and local search, as is often done in ACO algorithms. Section VI concludes the paper by briefly summarizing the obtained results.

II. THE PROBLEM AND THE ALGORITHM

Consider a minimization problem¹ (\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 $f(s)$, and Ω is a set of constraints, which defines the set of feasible candidate solutions. The goal of the minimization problem is to find an optimal solution s^* , i.e., a feasible candidate solution of minimum cost.

The combinatorial optimization problem (\mathcal{S}, f, Ω) is mapped on a problem that can be characterized by the following.²

- 1) A finite set $\mathcal{C} = \{c_1, c_2, \dots, c_{N_C}\}$ of *components*.
- 2) A finite set \mathcal{X} of *states* of the problem, defined in terms of all possible sequences $x = \langle c_i, c_j, \dots, c_k, \dots \rangle$ over the elements of \mathcal{C} . The length of a sequence x , i.e., the number of components in the sequence, is expressed by $|x|$. The maximum length of a sequence is bounded by a positive constant $n < +\infty$.
- 3) The set of (candidate) solutions \mathcal{S} is a subset of \mathcal{X} (i.e., $\mathcal{S} \subseteq \mathcal{X}$).
- 4) A set of feasible states $\tilde{\mathcal{X}}$, with $\tilde{\mathcal{X}} \subseteq \mathcal{X}$, defined via a problem-dependent test that verifies that it is not impos-

¹The obvious changes must be done if a maximization problem is considered.

²How this mapping can be done in practice has been described in a number of earlier papers on the ACO metaheuristic [5], [6].

sible to complete a sequence $x \in \tilde{\mathcal{X}}$ into a solution satisfying the constraints Ω .³

- 5) A nonempty set \mathcal{S}^* of optimal solutions, with $\mathcal{S}^* \subseteq \tilde{\mathcal{X}}$ and $\mathcal{S}^* \subseteq \mathcal{S}$.

Given the above formulation, artificial ants build candidate solutions by performing randomized walks on the completely connected, weighted graph $\mathcal{G} = (\mathcal{C}, \mathcal{L}, \mathcal{T})$, where the vertices are the components \mathcal{C} , the set \mathcal{L} fully connects the components \mathcal{C} , and \mathcal{T} is a vector gathering so-called pheromone trails τ .⁴ The graph \mathcal{G} is called construction graph.

Each artificial ant is put on a randomly chosen vertex of the graph and then it performs a randomized walk by moving at each step from vertex to vertex in the graph in such a way that the next vertex is chosen stochastically according to the strength of the pheromone currently on the arcs. While moving from one vertex to another of the graph \mathcal{G} , constraints Ω are used to prevent ants from building infeasible solutions. Once the ants have completed their walk, pheromone trails are updated. Formally, the solution construction behavior of a generic ant can be described as follows.

ANT_SOLUTION_CONSTRUCTION

While ($x_k \in \tilde{\mathcal{X}}$ and $x_k \notin \mathcal{S}$) do:

at each step k after building the sequence $x_k = \langle c_1, c_2, \dots, c_k \rangle$, select the next vertex (component) c_{k+1} randomly following

$$P(c_{k+1} = c \mid \mathcal{T}, x_k) = \begin{cases} \frac{\tau(c_k, c)^\alpha}{\sum_{\substack{y \in \mathcal{C} \\ (c_k, y) \in J_{c_k}}} \tau(c_k, y)^\alpha}, & \text{if } (c_k, c) \in J_{c_k} \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

where $0 < \alpha < +\infty$ is a parameter and a connection (c_k, y) belongs to J_{c_k} iff the sequence $x_{k+1} = \langle c_1, c_2, \dots, c_k, y \rangle$ is such that $x_{k+1} \in \tilde{\mathcal{X}}$. If at some point in the solution construction the set J_{c_k} is empty, the ant is dropped and its solution construction is terminated.

Once all the ants have terminated their ANT_SOLUTION_CONSTRUCTION procedure, a pheromone update phase is started in which pheromone trails are modified. Let \hat{s} be the best feasible solution found so far and s_t be the best feasible solution in the current algorithm iteration t ; $f(\hat{s})$ and $f(s_t)$ are the corresponding objective function values. The pheromone update procedure decreases by a small factor ρ , called the evaporation rate, the value of the pheromone trails on all the connections in \mathcal{L} and then increases the value of the pheromone trails on the connections belonging to \hat{s} (adding pheromone only to those edges that belong to the best-so-far

³By this definition the feasibility of a state $x \in \tilde{\mathcal{X}}$ should be interpreted in a weak sense. In fact, it does not guarantee that a completion s of x exists such that $s \in \tilde{\mathcal{X}}$.

⁴Pheromone trails can be associated to components, connections, or both. In the following, we will restrict our attention to the case in which pheromone trails are associated to connections, so that $\tau(i, j)$ is the pheromone associated to the connection between components i and j . It is straightforward to extend algorithms and proofs to the other cases.

solution is known in the literature as the global best offline pheromone update [5], [7] or global best for short in the following).

The pheromone update procedure is described as follows.

PHEROMONE_UPDATE

$\forall (i, j) : \tau(i, j) \leftarrow (1 - \rho) \cdot \tau(i, j)$

If $f(s_t) < f(\hat{s})$, then $\hat{s} \leftarrow s_t$

$\forall (i, j) \in \hat{s} : \tau(i, j) \leftarrow \tau(i, j) + g(\hat{s})$

$\forall (i, j) : \tau(i, j) \leftarrow \max\{\tau_{\min}, \tau(i, j)\}$

where ρ , $0 < \rho < 1$ is the evaporation rate, $\tau_{\min} > 0$ is a parameter, and $g(s)$, $0 < g(s) < +\infty$ is a function with $g : \mathcal{S} \mapsto \mathbb{R}^+$, $f(s) < f(s') \implies g(s) \geq g(s')$.

The algorithm is initialized as follows.

INITIALIZE

Generate a feasible solution s' and set $\hat{s} = s'$

$\forall (i, j)$, set $\tau(i, j) = \tau_0$

For each ant

select a start vertex c_1 according to some problem dependent criterion

set $k = 1$ and $x_k = \langle c_1 \rangle$

where τ_0 , $\tau_{\min} \leq \tau_0 < +\infty$, is a parameter.

After the initialization, the algorithm iterates through the procedures ANT_SOLUTION_CONSTRUCTION and PHEROMONE_UPDATE, until some termination condition is met. In the following, we will call this ACO algorithm $\text{ACO}_{\text{gb}, \tau_{\min}}$, where *gb* indicates that the global best pheromone update rule is used, while τ_{\min} indicates that a lower limit on the range of feasible pheromone trails is enforced. For the following we assume that $\tau_{\min} < g(s^*)$, which can be achieved by setting, e.g., $\tau_0 = g(s')/2$, where s' is the solution used to initialize $\text{ACO}_{\text{gb}, \tau_{\min}}$.

III. CONVERGENCE PROOF

We prove two theorems for the algorithm proposed in the previous section. First, we show that $\text{ACO}_{\text{gb}, \tau_{\min}}$ is guaranteed to find an optimal solution with a probability that can be made arbitrarily close to one if given enough time. Second, we prove that, after a fixed number of iterations t_0 has elapsed since the optimal solution was first found, the pheromone trails on the connections of the optimal solution are larger than those on any other connection. This result is then extended to show that an optimal solution can be constructed with a probability larger than $1 - \hat{\epsilon}(\tau_{\min}, \tau_{\max})$, where τ_{\max} is the maximum value the pheromones may take.

Before proving the first theorem, it is convenient to show that, due to pheromone evaporation, the maximum possible pheromone level τ_{\max} is bounded asymptotically.

Proposition 1: For any τ_{ij} , the following holds⁵:

$$\lim_{t \rightarrow \infty} \tau_{ij}(t) \leq \tau_{\max} = \frac{1}{\rho} \cdot g(s^*). \quad (2)$$

⁵In the proofs we write τ_{ij} instead of using $\tau(i, j)$ to ease notation.

Proof: The maximum possible amount of pheromone added to any edge (i, j) after any iteration is $g(s^*)$. Clearly, at iteration 1, the maximum possible pheromone trail is $(1 - \rho) \cdot \tau_0 + g(s^*)$, at iteration 2, it is $(1 - \rho)^2 \cdot \tau_0 + (1 - \rho) \cdot g(s^*) + g(s^*)$, etc. Hence, due to pheromone evaporation, the pheromone trail at iteration t is bounded by

$$\tau_{ij}^{\max}(t) = (1 - \rho)^t \cdot \tau_0 + \sum_{i=1}^t (1 - \rho)^{t-i} \cdot g(s^*).$$

Asymptotically, as $0 < \rho < 1$, this sum converges to

$$\tau_{\max} = \frac{1}{\rho} \cdot g(s^*)$$

Proposition 2: After an optimal solution has been found, remembering that $\tau_{ij}^*(t) \geq \tau_{\min}$ and that the global best pheromone update rule is used, we have that $\tau_{ij}^*(t)$ monotonically increases. The proof

$$\forall (i, j) \in s^* : \lim_{t \rightarrow \infty} \tau_{ij}^*(t) = \tau_{\max} = \frac{1}{\rho} \cdot g(s^*)$$

where τ_{ij}^* is the pheromone trail value on connections $(i, j) \in s^*$.

Proof: After an optimal solution has been found, remembering that $\tau_{ij}^*(t) \geq \tau_{\min}$ and that the global best pheromone update rule is used, we have that $\tau_{ij}^*(t)$ monotonically increases. The proof of Proposition 2 is basically a repetition of the proof of Proposition 1, restricted to the connections of the optimal solution (τ_0 is replaced by $\tau_{ij}^*(t^*)$ in the proof of Proposition 1, where t^* is the iteration when the first optimal solution has been found).

Proposition 1 says that for the following proof of Theorem 1, the only essential point is that $\tau_{\min} > 0$ because τ_{\max} will be bounded by pheromone evaporation. Proposition 2 states additionally that the pheromone trails on all connections of s^* , once an optimal solution has been found, converge to $\tau_{\max} = g(s^*)/\rho$.

We can now prove the following theorem.

Theorem 1: Let $P^*(t)$ be the probability that the algorithm finds an optimal solution at least once within the first t iterations. Then, for an arbitrary choice of a small $\epsilon > 0$ and for a sufficiently large t , it holds that

$$P^*(t) \geq 1 - \epsilon$$

and asymptotically

$$\lim_{t \rightarrow \infty} P^*(t) = 1.$$

Proof: Due to the pheromone trail limits τ_{\min} and τ_{\max} , we can guarantee that any feasible choice in (1) is done with a probability $p_{\min} > 0$. A trivial lower bound for p_{\min} can be given as⁶

$$p_{\min} \geq \hat{p}_{\min} = \frac{\tau_{\min}^{\alpha}}{(N_C - 1) \cdot \tau_{\max}^{\alpha} + \tau_{\min}^{\alpha}} \quad (3)$$

⁶For the derivation of this bound we consider the following "worst case" situation: the pheromone trail associated with the desired decision is τ_{\min} , while all the other feasible choices (there are at most $N_C - 1$) have an associated pheromone trail of τ_{\max} .

where N_C is the cardinality of the set \mathcal{C} of components.⁷ Then, any generic solution s' , including any optimal solution $s^* \in \mathcal{S}^*$, can be generated with a probability $\hat{p} \geq \hat{p}_{\min}^n > 0$, where $n < +\infty$ is the maximum length of a sequence. Because it is enough that one ant finds an optimal solution, a lower bound for $P^*(t)$ is given by

$$\hat{P}^*(t) = 1 - (1 - \hat{p})^t.$$

By choosing t sufficiently large, this probability can be made larger than any value $1 - \epsilon$ because we have that $\lim_{t \rightarrow \infty} \hat{P}^*(t) = 1$. ■

Theorem 2: Let t^* be the iteration when the first optimal solution has been found. Then a value t_0 exists such that the following holds:

$$\tau_{ij}(t) > \tau_{kl}(t)$$

$\forall (i, j) \in s^*, \forall (k, l) \in \mathcal{L} \wedge (k, l) \notin s^*, \text{ and } \forall t > t^* + t_0 = t^* + \lceil (1 - \rho)/\rho \rceil$

Proof: After a transition period t_0 has elapsed from the iteration in which the first optimal solution was found (i.e., for $t > t^* + t_0$), the pheromone trail on the connections used in the optimal solution is larger than that on any other feasible connection. In fact, due to the use of the global best pheromone update rule, only connections belonging to s^* increase their pheromone trails, while the pheromone trails of all other connections decrease by a factor ρ after each iteration until reaching the lower bound τ_{\min} .

We now give a bound on the length of the transition period t_0 . To do so, we assume the following worst case situation. Let (i, j) be a connection belonging to s^* with an associated pheromone trail at iteration t^* of $\tau_{ij}^*(t^*) = \tau_{\min}$. Also, let (k, l) be a connection not belonging to s^* with an associated pheromone trail at iteration t^* of $\tau_{kl}(t^*) = \tau_{\max}$. At iteration $t^* + t'$, $\tau_{ij}^*(t)$ becomes

$$\begin{aligned} & \tau_{ij}^*(t^* + t') \\ &= (1 - \rho)^{t'} \cdot \tau_{\min} + \sum_{i=0}^{t'-1} (1 - \rho)^i \cdot g(s^*) \\ &> t' \cdot (1 - \rho)^{(t'-1)} \cdot g(s^*) \end{aligned} \quad (4)$$

while the value of $\tau_{kl}(t)$ at iteration $t^* + t'$ is

$$\tau_{kl}(t^* + t') = \max\{\tau_{\min}, (1 - \rho)^{t'} \cdot \tau_{\max}\}.$$

For our purposes, the interesting case is when this maximum corresponds to $(1 - \rho)^{t'} \cdot \tau_{\max}$. Then we have that $\tau_{ij}^*(t^* + t') > \tau_{kl}(t^* + t')$ when

⁷It is easy to find tighter bounds for p_{\min} . One such bound can be derived if we consider the fact that, due to pheromone evaporation, no situation can arise in which $N_C - 1$ connections have a pheromone trail of τ_{\max} . In fact, if a connection with pheromone trail value τ_{\max} does not receive additional pheromone in the pheromone updates during i iterations, its pheromone level decreases to $(1 - \rho)^i \cdot \tau_{\max}$. Taking into account this effect, we can derive a tighter bound on p_{\min} as $\hat{p}'_{\min} = \tau_{\min}^{\alpha} / (\tau_{\min}^{\alpha} + \sum_{i=0}^{N_C-2} (1 - \rho)^{i\alpha} \cdot \tau_{\max}^{\alpha})$. In fact, this bound holds for any iteration $t > N_C$.

$$t' \cdot (1 - \rho)^{(t'-1)} \cdot g(s^*) > (1 - \rho)^{t'} \cdot \tau_{\max}$$

which is the case when

$$t' > \left\lceil \frac{\tau_{\max} \cdot (1 - \rho)}{g(s^*)} \right\rceil = \left\lceil \frac{(1 - \rho)}{\rho} \right\rceil \equiv t_0. \quad \blacksquare$$

From Theorem 2 it is immediate to see that, for any iteration $t > t^* + t_0$, any ant will be able to construct s^* by making, at each construction step, deterministically the choice associated with the largest pheromone trail.

In the limiting case, once the optimal solution has been found, we can give a lower bound estimate for an ant's probability of constructing an optimal solution when following the stochastic policy of the algorithm. Before proving this, we show in Proposition 3 that the pheromone trail of connections that are not in the optimal solution converge in the limit to the value τ_{\min} .

Proposition 3: Once an optimal solution has been found and for any $\tau_{ij}(t)$ such that $(i, j) \notin s^*$, it holds that

$$\lim_{t \rightarrow \infty} \tau_{ij}(t) = \tau_{\min}.$$

Proof: Because after the optimal solution has been found connections not belonging to the optimal solution do not receive any pheromone, their value can only decrease. In particular, after one iteration, $\tau_{ij}(t^* + 1) = \max\{\tau_{\min}, (1 - \rho) \cdot \tau_{\max}\}$ and after t' iterations, $\tau_{ij}(t^* + t') = \max\{\tau_{\min}, (1 - \rho)^{t'} \cdot \tau_{\max}\}$. It is then clear that for $t \rightarrow \infty$, $\tau_{ij}(t) \rightarrow \tau_{\min}$. \blacksquare

In fact, Proposition 3 can be made stronger by showing that the value τ_{\min} is reached in a finite number of iterations.

Proposition 4: Starting from iteration $t' \geq t^* + t_0$, it holds that

$$\forall (i, j) \notin s^*, \tau_{ij}(t) = \tau_{\min}$$

where t^* is the iteration when the first optimal solution s^* has been found and $t_0 = \lceil (\ln \tau_{\min} - \ln \tau_{\max}) / \ln(1 - \rho) \rceil$.

Proof: We can give a bound on t_0 by assuming that at iteration t^* for at least one $(i, j) \notin s^*$, it holds that $\tau_{ij}(t^*) = \tau_{\max}$. Following the same reasoning as in Proposition 3, we note that after t' iterations $\tau_{ij}(t^* + t') = \max\{\tau_{\min}, (1 - \rho)^{t'} \cdot \tau_{\max}\}$. Then, t_0 is the first iteration such that $(1 - \rho)^{t_0} \cdot \tau_{\max} \leq \tau_{\min}$. It is easy to show that this is the case for $t_0 = \lceil (\ln \tau_{\min} - \ln \tau_{\max}) / \ln(1 - \rho) \rceil$. \blacksquare

Corollary 1: Let t^* be the iteration when the first optimal solution has been found and $P(s^*, t, k)$ be the probability that an arbitrary ant k constructs s^* in the t th iteration, with $t > t^*$. Then it holds that

$$\lim_{t \rightarrow \infty} P(s^*, t, k) \geq 1 - \hat{\epsilon}(\tau_{\min}, \tau_{\max}).$$

Proof: Let ant k be located on component i and (i, j) be a connection of s^* . Remembering that ants choose the next component j in the set J_i , a lower bound $\hat{p}_{ij}^*(t)$ for the probability $p_{ij}^*(t)$ that ant k makes the ‘‘correct choice’’ (i, j) is given by the term

$$\hat{p}_{ij}^*(t) = \frac{(\tau_{ij}^*(t))^\alpha}{(\tau_{ij}^*(t))^\alpha + \sum_{(i,k) \notin s^*} (\tau_{ik}(t))^\alpha}.$$

Because of Propositions 2 and 3, we have

$$\begin{aligned} \hat{p}_{ij}^* &= \lim_{t \rightarrow \infty} \hat{p}_{ij}^*(t) \\ &= \frac{\lim_{t \rightarrow \infty} (\tau_{ij}^*(t))^\alpha}{\lim_{t \rightarrow \infty} \left[(\tau_{ij}^*(t))^\alpha + \sum_{(i,k) \notin s^*} (\tau_{ik}(t))^\alpha \right]} \\ &= \frac{\tau_{\max}^\alpha}{\tau_{\max}^\alpha + (N_C - 1) \cdot \tau_{\min}^\alpha}. \end{aligned} \quad (5)$$

Hence, in the limit, a lower bound for $P(s^*, t, k)$ is $\hat{p}_k^* = (\hat{p}_{ij}^*)^n$. Setting $\hat{\epsilon} = 1 - \hat{p}_k^*$ proves the corollary. \blacksquare

IV. DISCUSSION

In Section III, we proved two theorems about the convergence of $ACO_{\text{gb}, \tau_{\min}}$. In this section, we briefly discuss the meaning of these two theorems and we show how they relate to Gutjahr's previous convergence proof [14].

A. What Does the Proof Really Say?

It is instructive to understand what the proofs presented in Section III really tell us. Theorem 1 says that our algorithm does not rule out the possibility of finding the optimal solution, while Theorem 2 says that, once the optimal solution has been found, the pheromone trails on connections belonging to the optimal solution will become larger than those on any other connection. Based on this latter result, Corollary 1 gives a bound on the probability of constructing an optimal solution. On the other hand, the proofs do not say anything about the time required to find an optimal solution, which can be astronomically large. (A similar limitation applies to other well-known convergence proofs, such as those for simulated annealing [15], [17].)

It is interesting to stress the role of the strict inequality in the pheromone update rule of $ACO_{\text{gb}, \tau_{\min}}$ (see item 2 in the pheromone update rule of Section II). Suppose we replace the strict inequality with a better than or equal to inequality (i.e., we replace the symbol $<$ with the symbol \leq in the pheromone update rule). For a problem with several distinct global optima, this could lead to a situation in which we switch back and forth between distinct global optima. Although this does not affect Theorem 1, Theorem 2 would not hold anymore, i.e., we could not prove that all the pheromones settle on a single globally optimal solution. Interestingly, all those implementations of ACO algorithms that use the global best update rule also use a strict inequality, i.e., they update the global best solution only when an improved solution is found.

An important role in the proof of Theorem 1 is played by τ_{\min} and τ_{\max} : the smaller the ratio τ_{\max}/τ_{\min} , the larger the lower bound \hat{p}_{\min} given in that proof.⁸ This is important because the larger \hat{p}_{\min} , the smaller is the worst case estimate of the number of iterations t needed to assure that an optimal solution is found with a probability larger than $1 - \epsilon$. In fact, the tightest bound is obtained if all pheromone trails are the same, i.e., for the case of

⁸On the contrary, in Corollary 1, the larger the ratio τ_{\max}/τ_{\min} , the larger the asymptotic probability $P(s^*, t, k)$ that an ant k builds the optimal solution once it has been found.

uniformly random solution construction; in this case⁹ we would have $\hat{p}_{\min} = 1/N_C$. This, in a way, counterintuitive result is due to the fact that our proof is based on a worst case analysis: we need to consider the worst case situation in which the bias in the solution construction introduced by the pheromone trails is counterproductive and leads to suboptimal solutions,¹⁰ i.e., we have to assume that the pheromone trail level associated with the connection an ant needs to pass for constructing an optimal solution is τ_{\min} , while on the other connections, it is much higher—in the worst case, corresponding to τ_{\max} . As we said, however, the main contribution of the theorem, in practical terms, is to assure that if the algorithm runs long enough, it will find an optimal solution with a high probability. This result is important because, as we will see in Section V-A, the convergence proof for $\text{ACO}_{\text{gb},\tau_{\min}}$ can be extended to cover two of the experimentally best performing ACO algorithms.

B. Relationship to Gutjahr's Convergence Proof

Recently, Gutjahr [14] proposed a convergence proof for GBAS, an algorithm belonging to the ACO class. GBAS is very similar to $\text{ACO}_{\text{gb},\tau_{\min}}$ except that $\tau_{\min} = 0$ and the pheromone update rule changes the pheromones only when, in the current iteration, a solution at least as good as the best one found so far is generated. He proved the following theorem¹¹:

- 1) for each $\epsilon > 0$, for a fixed ρ , and for a sufficiently large number of ants, the probability P that a fixed ant constructs the optimal solution at iteration t is $P \geq 1 - \epsilon$ for all $t \geq t_0$, with $t_0 = t_0(\epsilon)$;
- 2) for each $\epsilon > 0$, for a fixed number of ants, and for an evaporation rate ρ sufficiently close to zero, the probability P that a fixed ant constructs the optimal solution at iteration t is $P \geq 1 - \epsilon$ for all $t \geq t_0$, with $t_0 = t_0(\epsilon)$.

There are a number of differences between Gutjahr's proof and ours, the most important concerning the type of convergence proved. In fact, in Theorem 1 we prove convergence in value (i.e., we prove that the algorithm will eventually find the optimal solution), while Gutjahr proves convergence in solution (i.e., he proves that the algorithm will converge to a situation in which it generates the optimal solution over and over). Gutjahr's proof is stronger than ours (it implies our result, but the reverse is not true), but our proof holds for any ACO algorithm as far as a lower bound $\tau_{\min} > 0$ and an upper bound $\tau_{\max} < +\infty$ to the pheromone trails exist. Therefore, while Gutjahr's proof holds only for GBAS, an ACO algorithm that has never been implemented and for which no experimental results are available, ours holds, as shown in Section V, for some of the best performing ACO algorithms published in the literature.

From a more technical point of view, the two proofs differ in a number of points that we summarize in the following.

⁹This fact is independent of the tightness of the lower bounds used in Theorem 1.

¹⁰In practice, however, as shown by the results of many published experimental works (see [5], [6], [10] for an overview) this does not happen, and the bias introduced by the pheromone trails does indeed help to speed up convergence to an optimal solution.

¹¹While finalizing this paper, Gutjahr [12] extended the convergence results of his earlier article [14] for two variants of GBAS, obtaining the very same convergence properties of simulated annealing [15], i.e., convergence of the current solution to an optimal solution with probability one.

- 1) Our Theorem 1 holds independent of the way pheromones are updated (to be exact, it holds, provided that $\tau_{\min} > 0$, for any pheromone update rule with $0 < \rho < 1$ and that does add a finite amount of pheromone trail), while Gutjahr's proof holds only for GBAS's particular pheromone update rule (in an extension of his theorem [13], he proves convergence under the condition that GBAS's pheromone update rule is applied at least in the final phases of the algorithm).
- 2) Our Theorems 1 and 2, as well as Corollary 1, hold for $\tau_{\min} > 0$, while in Gutjahr's proof, pheromone trails can go to zero.
- 3) Gutjahr proves that the probability of generating the optimal solution in each iteration goes to one as the number of iterations goes to infinity, while we can only prove, because of τ_{\min} , that it goes to $1 - \hat{\epsilon}$ (see Corollary 1).¹²
- 4) Our result is independent of the number of optimal solutions in \mathcal{S}^* , while one of the conditions for Gutjahr's theorem is that there is a single optimal solution (this limitation has been removed by Gutjahr in [13]).
- 5) In our Theorem 1, convergence is a function of τ_{\min} and ρ , while in Gutjahr's theorem, it is a function of the number of ants and of ρ .

V. ACO ALGORITHMS AND CONVERGENCE

As we already mentioned, from the point of view of the researcher interested in applications of the algorithm, the interesting part of our convergence proof is the one corresponding to Theorem 1, which states that $\text{ACO}_{\text{gb},\tau_{\min}}$ finds an optimal solution with arbitrarily large probability if run long enough.

It is, therefore, interesting that this theorem also applies to ACO algorithms that differ from $\text{ACO}_{\text{gb},\tau_{\min}}$ in the way the pheromone update procedure is implemented. In general, Theorem 1 applies to any ACO algorithm for which the probability $P(s)$ of constructing a solution $s \in \mathcal{S}$ always remains greater than a small constant $\epsilon > 0$. In $\text{ACO}_{\text{gb},\tau_{\min}}$, this is a direct consequence of the fact that $0 < \tau_{\min} < \tau_{\max} < +\infty$, which was obtained by: 1) explicitly setting a minimum value τ_{\min} for pheromone trails; 2) limiting the amount of pheromone that the ants may deposit after each iteration, i.e., $\forall s, g(s) < +\infty$; and 3) letting pheromone evaporate over time, i.e., by setting $\rho > 0$. We will call the class of ACO algorithms that satisfy these three conditions $\text{ACO}_{\tau_{\min}}$. $\text{ACO}_{\text{gb},\tau_{\min}}$ differs from $\text{ACO}_{\tau_{\min}}$ in that it additionally imposes the use of the global best pheromone update rule. Therefore, $\text{ACO}_{\text{gb},\tau_{\min}}$ can be seen as a particular case of $\text{ACO}_{\tau_{\min}}$. By definition, Theorem 1 holds for any algorithm in $\text{ACO}_{\tau_{\min}}$. On the contrary, this is not in general true for Theorem 2 and Corollary 1.

A. Algorithms in $\text{ACO}_{\tau_{\min}}$

In the following, we show that $\mathcal{M}\mathcal{M}\mathcal{A}\mathcal{S}$ and $\mathcal{A}\mathcal{C}\mathcal{S}$, two of the experimentally most successful ACO algorithms, fall into the $\text{ACO}_{\tau_{\min}}$ class.

¹²It should be said that, for practical purposes, this part of the proof is not very important because in optimization we are interested in finding the optimal solution and not in continuing to generate it once found. In fact, in any implementation of an iterative procedure for combinatorial optimization, the best solution found so far can be kept in memory and used as output of the procedure.

1) *MAX-MIN Ant System*: *MMAS* is one of the best-performing ACO algorithms and it was applied successfully to problems such as the well-known traveling salesman problem (TSP) and the quadratic assignment problem (QAP) [18], [19]. It is easy to show that *MMAS* belongs to $\text{ACO}_{\tau_{\min}}$. In fact, there are only two minor differences between *MMAS*¹³ and $\text{ACO}_{\text{gb},\tau_{\min}}$. First, *MMAS* uses an explicit value for τ_{\max} instead of an implicit one as done in $\text{ACO}_{\text{gb},\tau_{\min}}$.¹⁴ The main reason for this choice is that *MMAS* occasionally reinitializes the pheromone trails to τ_{\max} ; we refer to [19] for more details. Second, *MMAS* uses a somewhat more general pheromone update rule than $\text{ACO}_{\text{gb},\tau_{\min}}$. Like in $\text{ACO}_{\text{gb},\tau_{\min}}$, in *MMAS*, only one solution is used to select the connections on which to add pheromone, but it allows to choose between the iteration best solution s_t and the global best solution \hat{s} . It is, therefore, clear that Theorem 1 holds for *MMAS*.

In *MMAS*, it was shown experimentally that a good strategy is to choose more and more often the global best solution for the pheromone update, until reaching a situation in which pheromone is added only to connections belonging to \hat{s} . In this case, it is easy to adapt Theorem 2 so that it also holds (we assume here that no pheromone reinitialization is applied). It suffices to compute the transition period t_0 starting from $t' = \max\{t^*, \hat{t}\}$, where \hat{t} is the iteration after which only the global best solution adds pheromone.¹⁵ A similar reasoning applies to the proofs of Propositions 2, 3, and 4, and, therefore, to Corollary 1.

2) *Ant Colony System*: ACS [7], another very successful ACO algorithm, also belongs to $\text{ACO}_{\tau_{\min}}$. Yet, this is not as immediate to see as for *MMAS*. Therefore, we first give some more details on ACS and then we show why Theorem 1 also applies to it.

ACS differs in three main points from $\text{ACO}_{\text{gb},\tau_{\min}}$. First, ACS uses the pseudorandom proportional action choice rule. At each construction step, an ant has two possible choices: either it deterministically chooses the connection with the largest pheromone trail value or it performs a biased exploration according to (1). The first choice is made with probability q_0 , the second one with probability $(1 - q_0)$, where $0 \leq q_0 < 1$ is a parameter. Second, ACS does not apply pheromone evaporation to all connections. The update rule used in ACS is as follows.¹⁶

ACS_OFFLINE_PHEROMONE_UPDATE

If $f(s_t) < f(\hat{s})$, then $\hat{s} \leftarrow s_t$

$\forall (i, j) \in \hat{s} : \tau(i, j) \leftarrow (1 - \rho) \cdot \tau(i, j) + \rho \cdot g(\hat{s})$

where ρ is the pheromone evaporation.

Third, each ant in ACS uses a local pheromone trail update rule that the ants apply immediately after having crossed a connection during solution construction.

¹³We base the description of *MMAS* on the version in [19].

¹⁴In fact, this is a very minor difference because *MMAS* uses as an estimate of τ_{\max} the upper pheromone trail limit defined by Proposition 1. This is done by adapting τ_{\max} each time a new improved solution is found using $g(\hat{s})$ instead of $g(s^*)$ in (2), leading to a dynamically changing value of $\tau_{\max}(t)$.

¹⁵That is, replace every occurrence of t^* in Theorem 2 with t' .

¹⁶As in $\text{ACO}_{\text{gb},\tau_{\min}}$, but differently from *MMAS*, ACS uses only \hat{s} in the pheromone update.

ACS_ONLINE_STEP_BY_STEP_PHEROMONE_UPDATE

$(c_k, c) \in x_{k+1} : \tau(c_k, c) \leftarrow (1 - \xi) \cdot \tau(c_k, c) + \xi \cdot \tau_0$

where $\xi, 0 < \xi < 1$, and τ_0 are two parameters.¹⁷

The effect of the local updating rule is to make a chosen connection less desirable for the following ants. It is convenient to remark that the two pheromone update rules used in ACS are of the form $a_{k+1} = (1 - \psi) \cdot a_k + \psi \cdot b$, for $k \geq 1$, where a_{k+1} and a_k are $\tau_{ij}(t+1)$ and $\tau_{ij}(t)$ and, respectively, $b = g(\hat{s})$, τ_0 and $\psi = \rho, \xi$. Then, we have

$$a_k = (1 - \psi)^k \cdot a_0 + b \cdot [1 - (1 - \psi)^k]$$

whose limit is b as $k \rightarrow \infty$. The sequence decreases for $a_0 > b$ (with maximum a_0) and increases for $a_0 < b$ (with maximum b).

Now the question is: How does the convergence result of $\text{ACO}_{\text{gb},\tau_{\min}}$ transfer to ACS? First, we observe that in ACS the maximum amount of pheromone is limited by $\tau_{\max} = \tau_{\max}^{\text{ACS}} = g(s^*)$ (this bound is obtained without considering the local pheromone update). Moreover, the parameter τ_0 in ACS corresponds to τ_{\min} in $\text{ACO}_{\text{gb},\tau_{\min}}$, i.e., no pheromone trail value can fall below τ_0 . This is the case because τ_0 is chosen in such a way that it is smaller than¹⁸ $g(\hat{s})$ and in this case the limit of the sequence a_k from above corresponds to τ_0 , giving a lower bound on the pheromone trail of any solution component (i, j) .

The second step is to show that any feasible solution can be constructed with a nonzero probability. The easiest way to see this is to rewrite the probability of making some fixed choice (i, j) in ACS. Let us assume that connection (i, j) does not have the largest pheromone trail associated. Then, the probability of choosing connection (i, j) can be calculated as the product of the probability of making a randomized choice, which is $1 - q_0$, and the probability of choosing connection (i, j) in this randomized choice. A bound for the latter is given by \hat{p}_{\min} in (3). Therefore, a lower bound for the probability of making any specific choice at any construction step is $(1 - q_0) \cdot \hat{p}_{\min}$ and Theorem 1 directly applies to ACS.

As far as Theorem 2 is concerned, it also applies to ACS, except that it is no longer possible to easily derive a deterministic bound on the length of the transition period t_0 . However, it should be noted that, once an optimal solution s^* has been found, because of the global best pheromone update rule, the connections belonging to s^* are the only ones that continue to receive pheromone and that may increase their pheromone trails, while the pheromone trails on all the other connections can only decrease, until the lower limit τ_0 is reached (this happens each time they are used by some ant due to the local pheromone update rule).

B. Additional Features of ACO Algorithms

Many ACO algorithms [5], [6] include some features that are not present in $\text{ACO}_{\text{gb},\tau_{\min}}$. The most important are the use of

¹⁷The value $\tau_0, 0 < \tau_0 < g(s^*)$, is a small constant value, which in ACS is also used to initialize the pheromones. It can easily be guaranteed that $\tau_0 < g(s^*)$, for example, by first generating a solution s' and then setting $\tau_0 = g(s')/2$.

¹⁸See footnote 17.

local search to improve the solutions constructed by the ants¹⁹ and the use of heuristic information in the choice of the next connection or component. In fact, these two features are also frequently used in \mathcal{MMAS} and ACS. Therefore, a natural question is how these two extensions affect the convergence proofs for $\text{ACO}_{\text{gb},\tau_{\min}}$ and, hence, also those for $\text{ACO}_{\tau_{\min}}$.

Let us first consider the additional use of local search. Local search tries to improve an ant's solution s by iteratively applying small, local changes to it. Typically, the best solution s' found by the local search is returned and used to update the pheromone trails. It is rather easy to see that the use of local search does not affect the convergence properties of $\text{ACO}_{\text{gb},\tau_{\min}}$ —they only refer to the way solutions are constructed and hold irrespectively of the use of a local search. Despite the fact that local search does not affect the theoretical convergence behavior of $\text{ACO}_{\text{gb},\tau_{\min}}$, it is known that in practice ACO algorithms often become more effective when applying local search [5].

A priori available information on the problem can be used to derive heuristic information that biases the probabilistic decisions taken by the ants. When incorporating such heuristic information into $\text{ACO}_{\text{gb},\tau_{\min}}$, (1) becomes

$$P(c_{k+1} = c \mid \mathcal{T}, x_k) = \begin{cases} \frac{\tau(c_k, c)^\alpha \cdot \eta(c_k, c)^\beta}{\sum_{\substack{y \in \mathcal{C} \\ (c_k, y) \in \mathcal{J}_{c_k}}} \tau(c_k, y)^\alpha \cdot \eta(c_k, y)^\beta}, & \text{if } (c_k, c) \in \mathcal{J}_{c_k} \\ 0, & \text{otherwise} \end{cases} \quad (6)$$

where $\eta(c_k, c)$ measures the heuristic desirability of adding solution component c , and β is a parameter. Theorem 1 is not affected by the heuristic information if we have $0 < \eta(i, j) < +\infty$ for each $(i, j) \in \mathcal{L}$ and $\beta < \infty$. In an extension to Theorem 2, we still can guarantee that the components of s^* are the most probable ones to be chosen, if at each construction step the product $\tau(c_k, c^*)^\alpha \cdot \eta(c_k, c^*)^\beta$, where c^* is the component that we have to choose to construct s^* , is maximal. In fact, this can be guaranteed if $r_\tau^\alpha > r_\eta^\beta$, where $r_\tau = \tau_{\max}/\tau_{\min}$ and $r_\eta = \eta_{\max}/\eta_{\min}$ with η_{\min} and η_{\max} being the smallest and largest possible heuristic information for any connection $(i, j) \in \mathcal{L}$. Under this condition, a misleading heuristic information is made up by a larger range of possible pheromone trails. If this condition is verified, we also can compute a bound for the length t_0 of the transition period as in Theorem 2, which then becomes $t_0^\eta > \lceil [(1 - \rho) \cdot \tau_{\max} \cdot \eta_{\max} / (g(s^*) \cdot \eta_{\min})] \rceil = \lceil r_\eta \cdot (1 - \rho) / \rho \rceil$.

As a final remark, we note that AS, a particularly important ACO algorithm because it is the ancestor of all ACO algorithms [3], [8], [9], as well as some of its variants (for example, elitist AS [3], [9] and the rank-based version of AS [2]) do not belong to $\text{ACO}_{\tau_{\min}}$. In fact, in these three algorithms, there is no lower bound to the value of pheromone trails that can therefore become null. It is interesting to note that ACS and \mathcal{MMAS} were

shown to perform better than AS and its variants on many standard benchmark problems such as the TSP and the QAP. Therefore, we are in the fortunate case in which ACO algorithms for which convergence can be proved theoretically also show better performance in practice.

VI. CONCLUSION

In this paper, we have proved two theorems that apply to the ACO algorithm called $\text{ACO}_{\text{gb},\tau_{\min}}$. The first theorem states that the probability of finding at least once an optimal solution P^* can be made greater than $1 - \epsilon$ for any small constant $\epsilon > 0$ if the algorithm is run for a sufficiently large number of iterations. We have then shown that Theorem 1 applies to a larger class of ACO algorithms called $\text{ACO}_{\tau_{\min}}$, which differ from $\text{ACO}_{\text{gb},\tau_{\min}}$ in that they can use any reasonable pheromone trail update rule.

The second theorem, which applies to $\text{ACO}_{\text{gb},\tau_{\min}}$, states that starting from a fixed number of iterations after the optimal solution has been found, the pheromone trails will be higher on the connections belonging to the optimal solution than on any other connection. Therefore, an ant that at each construction step chooses the connection with the highest pheromone trail will deterministically construct the optimal solution. Additionally, in Corollary 1, we proved that for $t \rightarrow \infty$ any fixed ant will produce the optimal solution during the t th iteration with probability $P \geq 1 - \hat{\epsilon}(\tau_{\min}, \tau_{\max})$, where τ_{\min} and τ_{\max} are the minimum and maximum values that can be taken by pheromone trails.

Finally, we have shown that some of these results can be extended to two of the most used and successful ACO algorithms, namely, \mathcal{MMAS} and ACS.

ACKNOWLEDGMENT

The authors would like to thank D. Fogel, three anonymous referees, and M. Birattari, C. Blum, M. Hardt, M. Middendorf, A. Roli, M. Sampels, and M. Zlochin for their helpful comments on drafts of this paper. They would also like to thank W. J. Gutjahr for spotting some errors in an earlier version and for some valuable comments that helped to improve the paper.

REFERENCES

- [1] M. Birattari, G. Di Caro, and M. Dorigo, "For a formal foundation of the ant programming approach to combinatorial optimization. Part 1: The problem, the representation, and the general solution strategy," ATR-Human Information Processing Laboratories, Kyoto, Japan, Tech Rep. TR-H-301, Dec. 2000.
- [2] B. Bullnheimer, R. F. Hartl, and C. Strauss, "A new rank-based version of the ant system: A computational study," *Central Eur. J. Oper. Res. Econom.*, vol. 7, no. 1, pp. 25–38, 1999.
- [3] M. Dorigo, "Optimization, learning and natural algorithms," Ph.D. dissertation (in Italian), Dipartimento di Elettronica, Politecnico di Milano, Milano, Italy, 1992.
- [4] M. Dorigo, E. Bonabeau, and G. Theraulaz, "Ant algorithms and stigmergy," *Future Gener. Comput. Syst.*, vol. 16, no. 8, pp. 851–871, 2000.
- [5] M. Dorigo and G. Di Caro, "The ant colony optimization meta-heuristic," in *New Ideas in Optimization*, D. Corne, M. Dorigo, and F. Glover, Eds. London, U.K.: McGraw-Hill, 1999, pp. 11–32.
- [6] M. Dorigo, G. Di Caro, and L. M. Gambardella, "Ant algorithms for discrete optimization," *Artif. Life*, vol. 5, no. 2, pp. 137–172, 1999.
- [7] M. Dorigo and L. M. Gambardella, "Ant colony system: A cooperative learning approach to the traveling salesman problem," *IEEE Trans. Evol. Comput.*, vol. 1, pp. 53–66, Apr. 1997.

¹⁹The ACO metaheuristic is also applicable to time-varying problems in which the topology and costs can change while solutions are built. In this paper, we consider only applications to static problems for which topology and costs remain fixed; in fact, the convergence proof presented in this paper is meaningless in the case of time-varying problems where an algorithm must be able to follow the dynamics inherent to the problem.

- [8] M. Dorigo, V. Maniezzo, and A. Colomi, "Positive feedback as a search strategy," Dipartimento di Elettronica, Politecnico di Milano, Milano, Italy, Tech. Rep. 91-016, 1991.
- [9] —, "Ant system: Optimization by a colony of cooperating agents," *IEEE Trans. Syst. Man Cybern. B*, vol. 26, pp. 29–41, Feb. 1996.
- [10] M. Dorigo and T. Stützle, "The ant colony optimization metaheuristic: Algorithms, applications and advances," in *Handbook of Mathematics*, F. Glover and G. Kochenberger, Eds. Norwell, MA: Kluwer, to be published.
- [11] P. P. Grassé, "La reconstruction du nid et les coordinations interindividuelles chez *bellicositermes natalensis* et *cubitermes* sp. La théorie de la stigmergie: essai d'interprétation du comportement des termites constructeurs," *Insectes Sociaux*, vol. 6, pp. 41–81, 1959.
- [12] W. J. Gutjahr, "ACO algorithms with guaranteed convergence to the optimal solution," *Info. Processing Lett.*, vol. 82, no. 3, pp. 145–153, 2002.
- [13] —, "A generalized convergence result for the graph-based ant system metaheuristic," Dept. Statistics and Decision Support Systems, Univ. Vienna, Vienna, Austria, Tech. Rep. 99-09, 1999.
- [14] —, "A graph-based ant system and its convergence," *Future Gener. Comput. Syst.*, vol. 16, no. 8, pp. 873–888, 2000.
- [15] B. Hajek, "Cooling schedules for optimal annealing," *Math. Oper. Res.*, vol. 13, no. 2, pp. 311–329, May 1988.
- [16] N. Meuleau and M. Dorigo, "Ant colony optimization and stochastic gradient descent," *Artif. Life*, vol. 8, no. 2, pp. 103–121, 2002.
- [17] F. Romeo and A. Sangiovanni-Vincentelli, "A theoretical framework for simulated annealing," *Algorithmica*, vol. 6, no. 3, pp. 302–345, May 1991.
- [18] T. Stützle and H. H. Hoos, "The $MA\mathcal{X}$ - MIN ant system and local search for the traveling salesman problem," in *Proceedings of the 1997 IEEE International Conference on Evolutionary Computation (ICEC'97)*, T. Bäck, Z. Michalewicz, and X. Yao, Eds. Piscataway, NJ: IEEE Press, 1997, pp. 309–314.
- [19] —, " $MA\mathcal{X}$ - MIN ant system," *Future Gener. Comput. Syst.*, vol. 16, no. 8, pp. 889–914, 2000.



Thomas Stützle received the Dipl. degree in business engineering from the University of Karlsruhe, Karlsruhe, Germany, in 1994 and the Ph.D. degree in computer science from Darmstadt University of Technology, Darmstadt, Germany, in 1998.

From 1998 to February 2000, he was a Marie Curie Fellow at IRIDIA, Université Libre de Bruxelles, Brussels, Belgium. He is currently an Assistant Professor with the Computer Science Department, Darmstadt University of Technology. His current research interests include stochastic local search, the empirical analysis of stochastic algorithms, and the search space analysis of combinatorial problems.



Marco Dorigo (S'92–M'93–SM'96) received the Laurea (Master of Technology) degree in industrial technologies engineering and the doctoral degree in information and systems electronic engineering from Politecnico di Milano, Milan, Italy, in 1986 and 1992, respectively, and the Agrégé de l'Enseignement Supérieur title from the Université Libre de Bruxelles, Brussels, Belgium, in 1995.

From 1992 to 1993, he was a Research Fellow with the International Computer Science Institute, Berkeley, CA. In 1993, he was a NATO-CNR Fellow and, from 1994 to 1996, a Marie Curie Fellow. Since 1996, he has been a Tenured Researcher of the Belgian National Fund for Scientific Research. His current research interests include ant algorithms and swarm intelligence, evolutionary computation, distributed models of computation, and reinforcement learning. In addition, he is interested in the application of these computational paradigms to autonomous robotics, combinatorial optimization, and telecommunications networks problems.

Dr. Dorigo is an Associate Editor of the IEEE TRANSACTIONS ON EVOLUTIONARY COMPUTATION and IEEE TRANSACTIONS ON SYSTEMS, MAN, AND CYBERNETICS. He is a Member of the Editorial Board of numerous international journals, including *Evolutionary Computation*, *Artificial Life*, *Adaptive Behavior*, *Journal of Heuristics*, *Cognitive Systems Research*, and *Journal of Genetic Programming and Evolvable Machines*.