ACO Algorithms with Guaranteed Convergence to the Optimal Solution

Walter J. Gutjahr
Department of Statistics and Decision Support Systems
University of Vienna

Keywords: Analysis of algorithms, Ant Colony Optimization, combinatorial problems, convergence, heuristic optimisation, randomized algorithms.

1 Introduction

Ant Colony Optimization (short: ACO) algorithms for the heuristic solution of combinatorial optimization problems enjoy a rapidly growing popularity (see, e.g., the survey article by Dorigo, Di Caro and Gambardella [4]). Formulating the approach explicitly as a metaheuristic, the so–called ACO metaheuristic, Dorigo and Di Caro [5] have emphasized its broad range of applicability. The starting point for ACO algorithms is a biological metaphor: Natural ant colonies are able to find shortest paths between ant–hill and food by a specific type of reinforcement learning, using local pheromone trails for information exchange. In [6] and [3], Dorigo, Maniezzo and Colorni have transferred this principle to the algorithmic solution of optimization problems. In the last decade, it has been recognized that not only routing problems as the famous traveling salesperson problem (TSP), but also any other type of combinatorial optimization problems can be encoded as "best path" problems and solved using the ant colony metaphor.

Intuitively, the basic idea of ACO is the following: In order to solve a best–path problem in a graph, random walks of a fixed number of "ants" through the graph are simulated. The transition probabilities of each ant are governed by two types of parameters assigned to the edges of the graph: (i) *pheromone* values representing, in some sense, the "common memory" of the ant colony (the "past" perspective), and (ii) *visibility* values computed by heuristic pre–evaluations of how promising a transition along each edge appears (the "future" perspective). The pheromone values are updated by a combination of two mechanisms: *evaporation*, a global reduction of the pheromone vectors by a certain factor per time unit simulating the process that in a natural environment, pheromone trails diminish in course of time, and *reinforcement* of pheromone on paths recognized as "good".

Although much experimental and empirical research has already been done on ACO algorithms, only few algorithmic variants of the ant metaheuristic have been subject to a theoretical *convergence analysis.* The two basic existing attempts in this direction are the following:

1. In [9] and [10], an ACO variant called GBAS (Graph–based Ant System) is analyzed. The main result is that the current solutions of GBAS tend, with a probability $\geq 1 - \epsilon$, to

the optimal solution of the given optimization problem. It is shown that ϵ can be made arbitrarily small by a suitable choice of either the number S of ants or of the so–called evaporation factor ρ. However, no bounds on S or ρ ensuring an ϵ of a particular size are presented, so the convergence remains "uncontrollable" in a certain sense: The user cannot be given a convergence guarantee with a pre–specified minimum probability, even if he/she is willing to adjust the parameters S and ρ of his/her system.

2. In a recent paper [14], Stützle and Dorigo analyze another ACO algorithm, a variant of the MAX–MIN Ant System developed by Stützle and Hoos [16]. They show that for this variant, the objective function value of the best solution found so far converges with probability one to the optimal objective function value. Nevertheless, the demonstrated convergence property is relatively weak, being also shared by Random Search. For the current *solutions*, Stützle and Dorigo show a lower bound for their limiting probability of being optimal, but by construction of the algorithm, the gap of this lower bound to unity cannot be completely closed.

The aim of the present paper is to overcome the limitations of both approaches. For the first time, a convergence property of the strength of that proved by Hajek [11] for Simulated Annealing is shown within the context of ACO: We are able to demonstrate for a particular ACO algorithm that its current solutions converge to an optimal solution with probability *exactly one*. More specifically, we show that by using suitable parameter schemes, it can be guaranteed that the optimal paths get *attractors* of the stochastic dynamic process realized by the algorithm. The ACO algorithm for which we are able to provide such a guarantee is a time–dependent modification of GBAS. It uses two alternative types of pheromone update. Depending on which of these two types is chosen, we call the GBAS modification:

1. GBAS/tdev (GBAS with time–dependent evaporation factor), resp.

2. GBAS/tdlb (GBAS with time–dependent lower pheromone bound).

Let us emphasize that GBAS/tdev and GBAS/tdlb are not problem-specific algorithms, but full metaheuristics, such that they can be applied in the overall range of (static) combinatorial optimization.

2 The Algorithm

As it is done in [14], we formulate the investigated algorithm without reference to *visibility* (attractiveness, desirability) values η_{kl} (for an explanation of this notion, the reader is referred to [4], [5] or [9]). The two variants GBAS/tdev and GBAS/tdlb of the algorithm differ by the used pheromone update rules.

Definition 2.1. Let an instance of a combinatorial optimization problem be given. By a *construction graph* for this instance, we understand a directed graph $\mathcal{C} = (\mathcal{V}, \mathcal{A})$ together with a function Φ with the following properties:

(1) In $\mathcal{C}$, a unique node is marked as the so–called *start node*.

(2) Let $\mathcal{W}$ be the set of (directed) paths w in $\mathcal{C}$ satisfying the following conditions:

(i) w starts at the start node of $\mathcal{C}$;

(ii) w contains each node of $\mathcal{C}$ at most once;

(iii) the last node on w has no successor node in $\mathcal{C}$ that is not already contained in w (i.e., w cannot be prolonged without violating (ii)).

Then Φ maps a subset $\bar{\mathcal{W}}$ of the set $\mathcal{W}$ onto the set of feasible solutions of the given problem instance. In other words: To each path w in $\bar{\mathcal{W}}$, there corresponds (via Φ) a feasible solution, and to each feasible solution, there corresponds (via Φ^{-1}) at least one path in $\bar{\mathcal{W}}$.

Based on Definition 2.1, the algorithm is the following:

Initialize pheromone trails τ_{kl} on the arcs (k, l) of the construction graph $\mathcal{C}$;
for iteration $n = 1, 2, \ldots$ {
  for ant $s = 1, \ldots, S$ {
    set k, the current position of the ant, equal to the start node of $\mathcal{C}$;
    set u, the current path of the ant, equal to the empty list;
    while (a feasible continuation (k, l) of the path u of the ant exists) {
      select successor node l with probability p_{kl}, where
        $p_{kl} = 0$, if continuation (k, l) is infeasible, and
        $p_{kl} = \tau_{kl} \,/\, (\sum_{(k,r)} \tau_{kr})$, where the sum is over all feasible (k, r), otherwise;
      continue the current path u of the ant to node l by adding arc (k, l) and setting $k := l$;
    }
  }
  do pheromone update;
}

Therein, *feasibility* of a continuation (k, l) of a current (partial) path u is defined by property (2) of Definition 2.1: The complete path must be an element of $\bar{\mathcal{W}}$, so the continuation $u + (k, l)$ obtained by concatenating u with arc (k, l) must be such that a path $w \in \bar{\mathcal{W}}$ with prefix $u + (k, l)$ exists. The objective function value assigned to a complete path w is always that of $\Phi(w)$, i.e., of the corresponding feasible solution.

Remark 1. Disregarding *visibility* helps to increase transparence of presentation and is not an essential restriction: The proofs can be extended to the case of non–constant visibility values along the same lines as in [9] or [10]. — In the algorithm, the nice idea by Stützle and Dorigo in [14] is borrowed to always reinforce the *first* found instead of the last found of equally good paths. Also this feature is not restrictive: If the user prefers an implementation where the last (the currently found) best solution is reinforced, our convergence proof can be extended to this case by using the arguments in [10].

Remark 2. Definition 2.1 is the definition of the construction graph from [9] in a slightly generalized form. The generalization concerns the restriction of feasible paths to a subset $\bar{\mathcal{W}}$, which could be achieved in [9] by other means, namely via the visibility values.[1]

[1] Theoretically, $\bar{\mathcal{W}}$ can be *any* subset of $\mathcal{W}$ (our results remain valid in this general case). For practical implementations, however, it is important to decide *quickly* whether a path $w \in \bar{\mathcal{W}}$ with prefix $u + (k, l)$ exists, so the choice of $\bar{\mathcal{W}}$ needs to be restricted to subsets where the decision on the membership of a path with given prefix is computationally easy.

Remark 3. In [9], it is shown that each static combinatorial optimization problem can be encoded as a problem of finding the best feasible path in a suitable construction graph. The simplest example is a TSP with n nodes, where the construction graph is a complete graph of order n, and to each feasible path, there corresponds a tour of the TSP in a straightforward way. For other problem types, other graph types can be used, and specific constraints on feasible paths can be represented by a restriction of $\mathcal{W}$ to some appropriate $\bar{\mathcal{W}}$. For some problem encodings, an artificial start node has to be introduced; often, also an artificial termination node (a sink of the graph) needs to be introduced in order to satisfy (iii) in property (2) of Definition 2.1.

Now let us specify the two algorithmic variants by describing their respective pheromone update rules. In both variants, the pheromone *initialization* is done by $\tau_{kl} = \tau_{kl}(1) = 1/|\mathcal{A}|$ for all arcs (k, l) of $\mathcal{C}$, where $|\mathcal{A}|$ is the number of arcs of $\mathcal{C}$. The common principle of the two pheromone update rules is that a certain share ρ (or ρ_n) of pheromone "evaporates" on all arcs, which is compensated afterwards by placing pheromone as a "reward" on the arcs of the best found path. We say then that this path has been reinforced in the current iteration.

Pheromone update for GBAS/tdev:

$$\tau_{kl}(n+1) := \begin{cases} (1-\rho_n)\,\tau_{kl}(n) + \rho_n \,/\, L(\hat{w}(n)), & \text{if } (k,l) \in \hat{w}(n), \\ (1-\rho_n)\,\tau_{kl}(n), & \text{otherwise,} \end{cases} \tag{1}$$

where $0 < \rho_n < 1$ $(n = 1, 2, \ldots)$. Therein, $\hat{w}(n)$ is the best path found until the end of iteration n, which is stored in a specific (list) variable and replaced each time some ant finds a path with lower cost (objective function) value. $L(w)$ is the length, i.e., the number of arcs, of path w.

The essential feature of this GBAS variant is that the evaporation factor ρ_n is chosen as *time–dependent*: its value is a function of the iteration index n. As in (ordinary) GBAS (see [9]), the pheromone trails are normalized in such a way that their sum is always unity. The reward $\rho_n \,/\, L(\hat{w}(n))$ in (1) is the particular increment achieving just this normalization.

Pheromone update for GBAS/tdlb:

$$\tau_{kl}(n+1) := \begin{cases} \max\left((1-\rho)\,\tau_{kl}(n) + \rho \,/\, L(\hat{w}(n)),\ \tau_{min}(n)\right), & \text{if } (k,l) \in \hat{w}(n), \\ \max\left((1-\rho)\,\tau_{kl}(n),\ \tau_{min}(n)\right), & \text{otherwise,} \end{cases} \tag{2}$$

where $0 < \rho < 1$, and $(\tau_{min}(n))$ is a sequence of real numbers. The remaining notation is as before. Now, the evaporation factor is constant (time–independent), but the pheromone trails are bounded below in a time–dependent way, i.e., by a function of the iteration index n. Contrary to GBAS/tdev, where the sum of pheromone trails is always unity, it may vary in GBAS/tdlb. Defining a lower pheromone bound is a basic idea of the so-called MAX-MIN Ant System developed by Stützle [13] and investigated in much detail by Stützle and Hoos [15], [16]. Already in these articles, the possibility of dynamically changing this bound is considered.

3 Description by Stochastic Processes

For the derivation of stochastic convergence properties, it is necessary to put the algorithmic variants of section 2 into an unambiguous probabilistic framework by assigning well–defined stochastic processes to them. This can be done by considering certain current state variables connected with the execution. For our results, it is sufficient to concentrate on

- the vector $\underline{\tau}(n)$ of pheromone values on the arcs of $\mathcal{C}$ at the beginning of iteration n, and
- the best found path $\hat{w}(n-1)$ at the beginning of iteration n, i.e., at the end of iteration $n-1$. ($\hat{w}(0)$ can be defined as an arbitrary fixed path.)

$\underline{\tau}(n)$ is a vector with $|\mathcal{A}|$ components. $\hat{w}(n-1)$ can be represented as a sequence of at most $|\mathcal{V}|$ node indices, $|\mathcal{V}|$ being the number of nodes. Since there are only finitely many of such sequences which may be arranged in, say, lexicographic order, $\hat{w}(n-1)$ can also be identified by a single integer index. Then, $(\underline{\tau}(n),\, \hat{w}(n-1))$ is an element of $\mathrm{IR}^{|\mathcal{A}|} \times \mathrm{IN}$. Note that for a deterministic sequence of such elements (i.e., a particular realization of the stochastic process), the ordinary mathematical convergence notion applies.

The following Lemma classifies the described stochastic process as a *Markov process* (see, e.g., [8]:

Lemma 3.1. Both for GBAS/tdev and GBAS/tdlb, the stochastic process with states

$$X_n = (\underline{\tau}(n),\, \hat{w}(n-1)) \quad (n = 1, 2, \ldots)$$

is an inhomogeneous Markov process in discrete time. (For inhomogeneous Markov processes, see, e.g., [12].)

Proof. We have to show the Markov property that the distribution of X_n only depends on X_{n-1}. Observe that the distribution of the paths $w^{(s)}(n-1)$ of the ants $s = 1, \ldots, S$ in iteration $n-1$ only depends on $\underline{\tau}(n-1)$. Furthermore, these paths, together with $\hat{w}(n-2)$, determine $\hat{w}(n-1)$, and together with $\hat{w}(n-1)$, they determine $\underline{\tau}(n)$ by the pheromone update rules. Hence the distribution of $(\underline{\tau}(n),\, \hat{w}(n-1))$ is completely determined by $(\underline{\tau}(n-1),\, \hat{w}(n-2))$. The Markov process is inhomogeneous since the pheromone update depends on n. □

As the proof of the Lemma shows, the transition function of the Markov process follows, for each of the two algorithmic variants, in a straightforward way from the description in section 2. Thus, the stochastic processes to which our results in the next section refer are well–defined. Let us remark that the process $(\underline{\tau}(n))$, without the second component, is *not* a Markov process and cannot be defined in a self–contained way.

4 Results

Theorem 4.1. Let, in the algorithm GBAS/tdev,

$$\rho_n \;\leq\; 1 - \frac{\log n}{\log(n+1)} \quad (n \geq N) \tag{3}$$

for some $N \geq 1$, and

$$\sum_{n=1}^{\infty} \rho_n = \infty. \tag{4}$$

Then, with probability one, the states $X_n = (\underline{\tau}(n), \hat{w}(n-1))$ of the assigned Markov process converge as $n \to \infty$ to one of the states $(\underline{\tau}[w^*], w^*)$ where w^* is an optimal path, and $\underline{\tau}[w^*]$ is defined by

$$\tau_{kl}[w^*] = \begin{cases} 1/L(w^*), & \text{if } (k,l) \in w^*, \\ 0, & \text{otherwise.} \end{cases} \tag{5}$$

In particular, if (3) and (4) are satisfied, there exists with probability one an optimal path w^* such that for a fixed ant s, its probability $P_{w^*}(n)$ of traversing path w^* in iteration n tends to unity, as $n \to \infty$.

Proof.
(a) By F_n, we denote the event that iteration n is the first iteration in which an optimal path is traversed by some ant. Consider a fixed optimal path w^*. We have

$$\neg F_1 \wedge \neg F_2 \wedge \ldots \;\Rightarrow\; w^* \text{ is never traversed,}$$

and hence (with P denoting the probability measure on the Markov process)

$$P(\neg F_1 \wedge \neg F_2 \wedge \ldots) \leq P(w^* \text{ is never traversed})$$

$$= \prod_{n=1}^{\infty} P(w^* \text{ not traversed in iteration } n \mid w^* \text{ not traversed in an iteration } i < n). \tag{6}$$

Now, for each fixed arc (k,l), the following lower bound on its pheromone trail at the end of iteration n can be derived: In the worst case, (k,l) has always only lost pheromone by evaporation and never obtained a reinforcement. W.l.o.g., let $N \geq 2$. Then, for $n \geq N$, by (3)

$$\tau_{kl}(n) = \left[\prod_{i=1}^{n-1}(1-\rho_i)\right] \tau_{kl}(1) \;\geq\; \left[\prod_{i=1}^{N-1}(1-\rho_i)\right] \left[\prod_{i=N}^{n-1} \frac{\log i}{\log(i+1)}\right] \tau_{kl}(1)$$

$$= \left[\prod_{i=1}^{N-1}(1-\rho_i)\right] \tau_{kl}(1) \cdot \frac{\log N}{\log n} \;=\; \frac{const}{\log n}. \tag{7}$$

Since there are less than $|\mathcal{V}|$ nodes available as feasible successors of a fixed node, and since all pheromone trails are less or equal 1, it follows that

$$p_{kl}(n) \;\geq\; \frac{const}{|\mathcal{V}| \cdot \log n} \quad (n \geq N).$$

Hence, the probability of a fixed ant s to traverse w^* in iteration n ($n \geq N$) is

$$\prod_{(k,l)\in w^*} p_{kl}(n) \;\geq\; \left(\frac{const}{|\mathcal{V}| \cdot \log n}\right)^{L(w^*)}.$$

Obviously, the lower bound on the r.h.s. is independent on what has happened before iteration n, such that the estimation also holds conditionally on arbitrary events in iteration $1, \ldots, n-1$. Therefore, an upper bound for the r.h.s. of (6) is

$$1 \cdot \prod_{n=N}^{\infty} \left[1 - \left(\frac{const}{|\mathcal{V}| \cdot \log n}\right)^{L(w^*)}\right]. \tag{8}$$

Taking the logarithm of this product, we obtain

$$\sum_{n=N}^{\infty} \log\left(1 - \left(\frac{const}{|\mathcal{V}| \cdot \log n}\right)^{L(w^*)}\right) \;\leq\; - \sum_{n=N}^{\infty} \left(\frac{const}{|\mathcal{V}| \cdot \log n}\right)^{L(w^*)} \;=\; -\infty,$$

since $\sum_n (\log n)^{-L}$ is a diverging series for each positive integer L. (Note that $(\log n)^L \leq const \cdot n$ for n sufficiently large.) As a consequence, the product (8) and therefore also the r.h.s. of (6) is zero, which proves that the event $F_1 \vee F_2 \vee \ldots$ that in some iteration an optimal path is traversed has probability one.

After the iteration where, for the first time, an optimal path has been traversed, $\hat{w}(n)$ is set equal to this path. So we can re–formulate the observation above as follows: With probability one, the stochastic process $X_n = (\underline{\tau}(n), \hat{w}(n-1))$ enters in some iteration into one of the sets $\mathcal{S}_{|\mathcal{A}|} \times \{w^*\}$, where w^* is (the index of) an optimal path, and $\mathcal{S}_p = \{(x_1, \ldots, x_p) \mid x_i \geq 0, \sum_{i=1}^{p} x_i = 1\}$ is the p-dimensional simplex.

(b) We show now that for each w^* taken from the set of optimal paths, the state $(\underline{\tau}[w^*], w^*)$ with $\underline{\tau}[w^*]$ given by (5) is a deterministic attractor of the Markov process (X_n) with attraction domain $\mathcal{S}_{|\mathcal{A}|} \times \{w^*\}$: If the realization $X_1(\omega), X_2(\omega), \ldots$ of the Markov process satisfies $X_n(\omega) \in \mathcal{S}_{|\mathcal{A}|} \times \{w^*\}$ for some n, then $\lim_{n\to\infty} X_n(\omega) = (\underline{\tau}[w^*], w^*)$. Together with part (a) of the proof, this obviously verifies the first assertion of the theorem.

Let m denote the index of the iteration where w^* is traversed for the first time, such that the process $(\underline{\tau}(n), \hat{w}(n-1))$ enters into the set $\mathcal{S}_{|\mathcal{A}|} \times \{w^*\}$ in iteration $m+1$. Then, in all iterations $n > m$, only path w^* is reinforced, and $\hat{w}(n) = w^*$. Using the upper row of the update rule (1), it can easily be verified by induction that for $(k,l) \in w^*$ and $r = 1, 2, \ldots$,

$$\tau_{kl}(m+r) \;=\; \left[\prod_{n=m}^{m+r-1} (1-\rho_n)\right] \tau_{kl}(m) \;+\; \frac{1}{L(w^*)} \sum_{i=0}^{r-1} \rho_{m+i} \prod_{j=i+1}^{r-1} (1-\rho_{m+j}). \tag{9}$$

(Therein, the convention $\prod_{j\in J} a_j = 1$ for $J = \emptyset$ is used.) Due to (4), the series $\sum \rho_n$ diverges, i.e., $\prod_{n=1}^{\infty}(1-\rho_n) = 0$. Therefore, in (9),

$$\lim_{r\to\infty} \left[\prod_{n=m}^{m+r-1} (1-\rho_n)\right] \tau_{kl}(m) \;=\; \tau_{kl}(m) \prod_{n=m}^{\infty} (1-\rho_n) \;=\; 0, \tag{10}$$

which shall be needed further below.

Now consider an arc $(k,l) \notin w^*$. After iteration m, it is never reinforced, so

$$\tau_{kl}(m+r) \;=\; \left[\prod_{n=m}^{m+r-1} (1-\rho_n)\right] \tau_{kl}(m) \;\to\; 0 \quad (r \to \infty) \tag{11}$$

as in (10). Hence also the *sum* of the pheromone trails on arcs $\notin w^*$ tends to zero, and so by the normalization property, the sum of the pheromone trails on arcs $\in w^*$ must tend to unity. By (9) and (10), however, $\limsup_{r\to\infty} \tau_{kl}(m+r)$ has the same value for each $(k,l) \in w^*$ (the second term on the r.h.s. of (9) does not depend on (k,l)), and the sum of these values is 1. The same holds for $\liminf_{r\to\infty} \tau_{kl}(m+r)$. Therefore, for $(k,l) \in w^*$,

$$\limsup_{r\to\infty} \tau_{kl}(m+r) \;=\; \liminf_{r\to\infty} \tau_{kl}(m+r) \;=\; \lim_{r\to\infty} \tau_{kl}(m+r) \;=\; 1 \,/\, L(w^*). \tag{12}$$

Together with the trivial convergence $\lim w^* = w^*$ of the second component of X_n, this yields $\lim_{n\to\infty} X_n = (\underline{\tau}[w^*], w^*)$.

(c) In particular, if the process has entered a set $\mathcal{S}_{|\mathcal{A}|} \times \{w^*\}$, it follows immediately from the limits (11) and (12) of the pheromone trails and the definition of the transition probability p_{kl} that $p_{kl}(n) \to 1 \quad \text{for } (k,l) \in w^*$. Taking the product over all $(k,l) \in w^*$ yields convergence of $P_{w^*}(n)$ to unity. Since an optimal path w^* satisfying the condition above exists with probability one, the second assertion of the theorem follows. □

Remark. For a formally correct interpretation of the second assertion of Theorem 4.1 it should be observed that the probability $P_{w^*}(n)$ of a traversal of w^* is a function of the pheromone trails $\tau_{kl}(n)$. In our Markov process framework, these trails are components of the states and hence random variables, so also the values $P_{w^*}(n)$ $(n = 1, 2, \ldots)$ form, for each fixed w^*, a sequence of random variables. For a fixed realization of the process, $P_{w^*}(n)$ $(n = 1, 2, \ldots)$ is a sequence of real numbers which may or may not converge. By letting w^* run through all optimal paths, we get (in general) several sequences of this type. The theorem states that the case where one of these sequences tends to unity occurs with probability one.

The following corollary indicates concrete schemes for choosing the evaporation factors in accordance with Theorem 4.1:

Corollary. Let, in the algorithm GBAS/tdev,

$$\rho_n = \frac{c_n}{n \log(n+1)} \quad (n \geq 1) \qquad \text{with } 0 < \lim_{n\to\infty} c_n < 1.$$

Then the assertions of Theorem 4.1 hold.

Proof. We show that the conditions (3) and (4) are satisfied for such schemes. As $n \to \infty$,

$$1 - \frac{\log n}{\log(n+1)} = \frac{\log\left(1 + \frac{1}{n}\right)}{\log(n+1)} \sim \frac{1}{n \log(n+1)},$$

where $a_n \sim b_n$ $(n \to \infty)$ means $a_n / b_n \to 1$ $(n \to \infty)$.

Thus, if $c_n \to c \in]0, 1[$,

$$\frac{c_n}{n \log(n+1)} < 1 - \frac{\log n}{\log(n+1)}$$

for sufficiently large n, which shows (3). On the other hand,

$$\sum_{n=1}^{\infty} \frac{1}{n \log(n+1)} \geq \int_1^{\infty} \frac{dx}{x \log(x+1)} = \infty,$$

and hence also (with a sufficiently large N)

$$\sum_{n=1}^{\infty} \frac{c_n}{n \log(n+1)} \geq \sum_{n=N}^{\infty} \frac{c/2}{n \log(n+1)}$$

diverges, which shows (4). □

Theorem 4.2. Let, in the algorithm GBAS/tdlb,

$$\tau_{min}(n) = \frac{c_n}{\log(n+1)} \quad (n \geq 1)$$

with $\lim_{n\to\infty} c_n > 0$. Then the assertions of Theorem 4.1 hold.

Proof. The proof is an analogue to that of Theorem 4.1. Let us only go into the points that are different. In part (a) of the proof of Theorem 4.1, a worst–case estimation for the pheromone trail on a never reinforced arc (k,l) has been performed. In the context of GBAS/tdlb, we can use the lower bound $\tau_{min}(n)$ for that purpose and obtain

$$\tau_{kl}(n) \geq \tau_{min}(n) = \frac{c_n}{\log(n+1)} \geq \frac{c}{2\,\log(n+1)}$$

for sufficiently large n, where $c = \lim_{n\to\infty} c_n$. This is essentially the same lower bound as in (7), and the subsequent arguments can be carried over in an analogous way. As to part (b), we must be aware that $\underline{\tau}(n)$ is now not restricted anymore to the simplex $\mathcal{S}_{|\mathcal{A}|}$, since the pheromone trails are not normalized to sum $= 1$. In particular, the fact that pheromone vanishes on arcs $(k,l) \notin w^*$ does not imply anymore that the sum of pheromone trails on arcs $(k,l) \in w^*$ tends to unity. However, by a slight extension of the arguments, it is possible to show that if the set $\mathbb{R}^{|\mathcal{A}|} \times \{w^*\}$ is entered by the process (X_n), where w^* is an optimal path, one has, as before, $\lim_{n\to\infty} X_n = (\underline{\tau}[w^*], w^*)$ with $\underline{\tau}[w^*]$ given by (5). This can be seen as follows: Let again m denote the index of the iteration where w^* is traversed for the first time. A lower bound on $\tau_{kl}(n)$ for $(k,l) \in w^*$ is obtained by neglecting the possibility that $\tau_{kl}(n+1)$ is set equal to $\tau_{min}(n)$. In this way, we obtain, setting $\rho_n \equiv \rho$ in (9), that

$$\tau_{kl}(m+r) \geq (1-\rho)^r\,\tau_{kl}(m) + \frac{\rho}{L}\sum_{i=0}^{r-1}(1-\rho)^{r-i-1} \to \frac{\rho}{L}\sum_{j=0}^{\infty}(1-\rho)^j = \frac{1}{L}. \quad (r\to\infty) \quad (13)$$

As a consequence, $\tau_{kl}(m+r) > 1/(2L)$ for sufficiently large r. On the other hand, since $\tau_{min}(n) \to 0$ $(n\to\infty)$, we have $\tau_{min}(m+r) < 1/(2L)$ for sufficiently large r, such that the update type of setting $\tau_{kl}(n+1)$ equal to $\tau_{min}(n)$ is not applied anymore for large r values. Hence, for r large enough, the "greater equal" in (13) is to be replaced by "equal", which means that for $(k,l) \in w^*$, we have $\lim_{r\to\infty} \tau_{kl}(m+r) = 1/L$. For $(k,l) \notin w^*$, it is easily seen that $\lim_{r\to\infty} \tau_{kl}(m+r) = 0$, because $\tau_{min}(n) \to 0$ $(n\to\infty)$. The remainder of the proof, including part (c), is analogous to that of Theorem 4.1. $\square$

Remark. We admit that at present, we are not able to give clues how to choose the parameter schemes in Theorems 4.1 and 4.2 to speed up convergence in a *best–possible* way. Within the range where the given conditions are satisfied, the algorithm keeps a suitable balance between the two contradicting aims of exploring the search space and favoring good solutions. We could outline "extreme" schemes within the indicated range that emphasize either the first or the second aim, but it is very difficult to determine the *optimal* point of this range in a theoretical analysis.

5 Conclusion

By the presented theorems, the results of [9], [10] and [14] have been sharpened for a specific ACO algorithm to the strength of the well–known convergence property of the Simulated Annealing metaheuristic. As in Simulated Annealing (cf. [1]), it turns out that a convergence guarantee can be obtained by a suitable speed of "cooling" (i.e., reduction of the influence of randomness): The geometric pheromone decrement on not reinforced arcs effected by constant evaporation factor is too fast and leads (in general) to premature convergence to suboptimal solutions. On the other hand, introducing a fixed lower pheromone bound stops cooling at some point and leads to random–search–like behavior without convergence. In between lies a compromise of allowing pheromone trails to tend to zero, but slower than geometrically. This can be achieved either by decreasing evaporation factors, or else by "slowly" decreasing lower pheromone bounds. In a certain window of the cooling speed, we get convergence to the optimal solution with probability one.

We do not affirm that the theoretical cooling speeds indicated by Theorems 4.1 – 4.2 are also the most *efficient* ones, where efficiency is measured by the average runtime required to find a solution of a sufficiently good quality (say, only $p\%$ worse than the best solution with some pre–defined p). In the typical area of application for ACO, the area of NP–complete combinatorial optimization problems, we cannot expect to obtain an algorithm providing optimal solutions in a short computation time. Again as in Simulated Annealing, it might turn out that faster cooling than indicated by the theoretical scheme is advantageous for finite–time (i.e., real) computing: For getting quick convergence, it may be worthwhile to pay the price of convergence to suboptimal solutions. However, experimental studies with *slightly* decreasing evaporation factors or lower pheromone bounds falling *slightly* slower than geometrically might be interesting, especially for applications where the user is willing to invest a high amount of computation time for obtaining excellent solution quality. In addition to the (so–called "elitist") pheromone update mechanisms investigated here (and in [14]), we also suggest computational experiments with decreasing evaporation factors and/or decreasing lower pheromone bounds for other empirically successful update mechanisms, such as the *rank–based* update rule introduced by Bullheimer, Hartl and Strauss [2]. It would not be a surprise if some moderate form of retarding the cooling process could, in a considerable number of cases, be able to further improve the performance of present ACO implementations.

Future theoretical research should deal with the speed of convergence, and address the question of the expected running time until the optimum is reached. Some interesting results concerning the last–mentioned question have recently been found for the field of Evolutionary Algorithms (see [7]). Also the Ant Colony Optimization approach might be accessible to such investigations.

References

[1] Aarts, E., Korst, J., *Simulated Annealing and Boltzmann Machines*, Wiley (1989).

[2] Bullnheimer, B., Hartl, R. F., Strauss, C., "A new rank–based version of the Ant System: A computational study", *Central European Journal for Operations Research* 7 (1) (1999), pp. 25–38.

[3] Dorigo, M., "Optimization, Learning and Natural Algorithms", PhD Thesis, Dept. of Electronics, Politecnico di Milano, Italy (1992).

[4] Dorigo, M., Di Caro, G., Gambardella, L.M., "Ant Algorithms for Discrete Optimization", *Artificial Life* 5(2) (1999), pp. 137–172.

[5] Dorigo, M., Di Caro, G., "The Ant Colony Optimization metaheuristic", in: *New Ideas in Optimization*, D. Corne, M. Dorigo, F. Glover (eds.), pp. 11-32, McGraw–Hill (1999).

[6] Dorigo, M., Maniezzo, V., Colorni, A., "The Ant System: An autocatalytic optimization process", Technical Report 91–016, Dept. of Electronics, Politecnico di Milano, Italy (1991).

[7] Droste, S., Jansen, T., Wegener, I., "On the analysis of the (1+1) evolutionary algorithm", to appear in *Theoretical Computer Science.*

[8] Gallager, R. G., *Discrete Stochastic Processes*, Kluwer (1996).

[9] Gutjahr, W.J., "A graph–based Ant System and its convergence", *Future Generation Computer Systems* 16 (2000), pp. 873–888.

[10] Gutjahr, W.J., "A generalized convergence result for the Graph–based Ant System", Technical Report 99-09, Dept. of Statistics and Decision Support Systems, University of Vienna, Austria (1999).

[11] Hajek, B., "Cooling schedules for optimal annealing", *Mathematics of OR*, 13 (1988), pp. 311–329.

[12] Seneta, E., *Non–negative Matrices and Markov Chains*, 2nd ed., Springer (1981).

[13] Stützle, T., "MAX-MIN Ant System for the quadratic assignment problem", Technical Report AIDA-97-4, FG Intellektik, TU Darmstadt, Germany (1997).

[14] Stützle, T., Dorigo, M., "A short convergence proof for a class of ACO algorithms", Technical Report IRIDIA/2000-35, Universite Libre de Bruxelles, Belgium.

[15] Stützle, T., Hoos, H.H., "The MAX-MIN Ant system and local search for the travelling salesman problem", in: T. Baeck, Z. Michalewicz and X. Yao (eds.), *Proc. ICEC '97* (Int. Conf. on Evolutionary Computation) (1997), pp. 309–314.

[16] Stützle, T., Hoos, H.H., "MAX-MIN Ant System", *Future Generation Computer Systems*, 16 (2000), pp. 889–914.