



# *Ant Colony Optimization*

*Marco Dorigo and Thomas Stützle*



# Ant Colony Optimization

Marco Dorigo  
Thomas Stützle



30807682

A Bradford Book  
The MIT Press  
Cambridge, Massachusetts  
London, England



© 2004 Massachusetts Institute of Technology

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

This book was set in Times New Roman on 3B2 by Asco Typesetters, Hong Kong. Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Dorigo, Marco.

Ant colony optimization / Marco Dorigo, Thomas Stützle.

p. cm.

“A Bradford book.”

Includes bibliographical references (p. ).

ISBN 978-0-262-04219-2(alk. paper)

1. Mathematical optimization. 2. Ants—Behavior—Mathematical models. I. Stützle, Thomas. II. Title.

QA402.5.D64 2004

519.6—dc22

2003066629

10 9 8 7 6 5



To Serena and Roberto

To Maria José and Alexander

## Preface

Ants exhibit complex social behaviors that have long since attracted the attention of human beings. Probably one of the most noticeable behaviors visible to us is the formation of so-called ant streets. When we were young, several of us may have stepped on such an ant *highway* or may have placed some obstacle in its way just to see how the ants would react to such disturbances. We may have also wondered where these ant highways lead to or even how they are formed. This type of question may become less urgent for most of us as we grow older and go to university, studying other subjects like computer science, mathematics, and so on. However, there are a considerable number of researchers, mainly biologists, who study the behavior of ants in detail.

One of the most surprising behavioral patterns exhibited by ants is the ability of certain ant species to find what computer scientists call shortest paths. Biologists have shown experimentally that this is possible by exploiting communication based only on pheromones, an odorous chemical substance that ants may deposit and smell. It is this behavioral pattern that inspired computer scientists to develop algorithms for the solution of optimization problems. The first attempts in this direction appeared in the early '90s and can be considered as rather “toy” demonstrations, though important for indicating the general validity of the approach. Since then, these and similar ideas have attracted a steadily increasing amount of research—and ant colony optimization (ACO) is one outcome of these research efforts. In fact, ACO algorithms are the most successful and widely recognized algorithmic techniques based on ant behaviors. Their success is evidenced by the extensive array of different problems to which they have been applied, and moreover by the fact that ACO algorithms are for many problems among the currently top-performing algorithms.

### Overview of the Book

This book introduces the rapidly growing field of ant colony optimization. It gives a broad overview of many aspects of ACO, ranging from a detailed description of the ideas underlying ACO, to the definition of how ACO can generally be applied to a wide range of combinatorial optimization problems, and describes many of the available ACO algorithms and their main applications. The book is divided into seven chapters and is organized as follows.

Chapter 1 explains how ants find shortest paths under controlled experimental conditions, and illustrates how the observation of this behavior has been translated into working optimization algorithms.

In chapter 2, the ACO metaheuristic is introduced and put into the general context of combinatorial optimization. Basic notions of complexity theory, such as  $\mathcal{NP}$ -hardness, are given and other major metaheuristics are briefly overviewed.

Chapter 3 is dedicated to the in-depth description of all the major ACO algorithms currently available in the literature. This description, which is developed using the traveling salesman problem as a running example, is completed by a guide to implementing the algorithms. A short description of a basic C implementation, as well as pointers to the public software available at [www.aco-metaheuristic.org/aco-code/](http://www.aco-metaheuristic.org/aco-code/), is given.

Chapter 4 reports on what is currently known about the theory of ACO algorithms. In particular, we prove convergence for a specific class of ACO algorithms and we discuss the formal relation between ACO and other methods such as stochastic gradient descent, mutual-information-maximizing input clustering, and cross-entropy.

Chapter 5 is a survey of current work exploiting ACO to solve a variety of combinatorial optimization problems. We cover applications to routing, assignment, scheduling, and subset problems, as well as a number of other problems in such diverse fields as machine learning and bioinformatics. We also give a few “application principles,” that is, criteria to be followed when attacking a new problem using ACO.

Chapter 6 is devoted to the detailed presentation of AntNet, an ACO algorithm especially designed for the network routing problem, that is, the problem of building and maintaining routing tables in a packet-switched telecommunication network.

Finally, chapter 7 summarizes the main achievements of the field and outlines some interesting directions for future research.

Each chapter of the book (with the exception of the last chapter) ends with the following three sections: bibliographical remarks, things to remember, and exercises.

- Bibliographical remarks, a kind of short annotated bibliography, contains pointers to further literature on the topics discussed in the chapter.
- Things to remember is a bulleted list of the important points discussed in the chapter.
- Exercises come in two forms, thought exercises and computer exercises, depending on the material presented in the chapter.

Finally, there is a long list of references about ACO algorithms that gives a lot of pointers to more in-depth literature.

Overall, this book can be read easily by anyone with a college-level scientific background. The use of mathematics is rather limited throughout, except for chapter 4, which requires some deeper knowledge of probability theory. However, we assume

that the reader is familiar with some basic notions of graph theory, programming, and probabilities. The book is intended primarily for (1) academic and industry researchers in operations research, artificial intelligence, and computational intelligence; (2) practitioners willing to learn how to implement ACO algorithms to solve combinatorial optimization problems; and (3) graduate and last-year undergraduate students in computer science, management science, operations research, and artificial intelligence.

## Acknowledgments

The field of ant colony optimization has been shaped by a number of people who have made valuable contributions to the development and success of the field.

First of all, we wish to acknowledge the contributions of Alberto Colorni and Vittorio Maniezzo. Alberto and Vittorio collaborated closely with Marco Dorigo in the definition of the first ACO algorithms while Marco was a doctoral student at Politecnico di Milano, in Milan, Italy. Without their contribution, there would probably be no ACO research to describe. Our thoughts turn next to Jean-Louis Deneubourg and Luca Maria Gambardella. Jean-Louis, a recognized expert in the study of social insects, provided the inspiration (as described in chapter 1 of this book) for the ACO work. Luca, a computer scientist with a strong feeling for practical applications, was the one who most helped in transforming ACO from a fascinating toy into a competitive metaheuristic.

More generally, many researchers have written papers on ACO (applications, theoretical results, and so on). This book is clearly influenced by their research and results, which are reported in chapter 5.

Several colleagues and students of ours have checked large parts of the book. We appreciated very much the comments by Maria Blesa, Christian Blum, Julia Handl, Elena Marchiori, Martin Middendorf, Michael Samples, and Tommaso Schiavinotto. In addition, we would like to thank those colleagues who checked parts of the book: Mauro Birattari, Roberto Cordone, Gianni Di Caro, Karl Dörner, Alex Freitas, Luca Maria Gambardella, Jose Antonio Gámez, Walther Gutjahr, Richard Hartl, Holger Hoos, Joshua Knowles, Guillermo Leguizamón, John Levine, Helena Lourenço, Max Manfrin, Vittorio Maniezzo, Daniel Merkle, José Miguel Puerta, Marc Reimann, Andrea Roli, Alena Shmygelska, Krzysztof Socha, Christine Solnon, and Mark Zlochin. Our special thanks goes to Cristina Versino, for providing the ant drawings used in figures 1.7 and 3.2, and to all the people at the IRIDIA and Intellectics groups, for providing a stimulating scientific and intellectual environment in which to work.

People at MIT Press, and in particular Robert Prior, have greatly helped to make this project successful. We thank all of them, and in particular Bob, for gently pressing us to deliver the draft of this book.

Final thanks go to our families, in particular to our wives Laura and Maria José, who have constantly provided the comfortable environment conducive to successfully completing this book, and to our children Luca, Alessandro, and Alexander, who give meaning to our lives.

Marco Dorigo acknowledges support from the Belgian FNRS, of which he is a senior research associate. The writing of this book has been indirectly supported by the numerous institutions who funded the research of the two authors through the



years. We wish to thank the Politecnico di Milano, Milan, Italy; the International Computer Science Institute, Berkeley, California; IDSIA, Lugano, Switzerland; the Intellectics Group at Darmstadt University of Technology, Germany; the IRIDIA group at the Université Libre de Bruxelles, Brussels, Belgium; and the Improving Human Potential programme of the CEC, who supported this work through grant HPRN-CT-1999-00106 to the Research Training Network “Metaheuristics Network.” 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.

# Contents

|  |            |
|--|------------|
| Preface  | ix         |
| Acknowledgments  | xiii       |
| <b>1 From Real to Artificial Ants</b>  | <b>1</b>   |
| 1.1 Ants' Foraging Behavior and Optimization                                   | 1          |
| 1.2 Toward Artificial Ants   | 7          |
| 1.3 Artificial Ants and Minimum Cost Paths                                     | 9          |
| 1.4 Bibliographical Remarks  | 21         |
| 1.5 Things to Remember   | 22         |
| 1.6 Thought and Computer Exercises   | 23         |
| <b>2 The Ant Colony Optimization Metaheuristic</b>                             | <b>25</b>  |
| 2.1 Combinatorial Optimization   | 25         |
| 2.2 The ACO Metaheuristic  | 33         |
| 2.3 How Do I Apply ACO?  | 38         |
| 2.4 Other Metaheuristics   | 46         |
| 2.5 Bibliographical Remarks  | 60         |
| 2.6 Things to Remember   | 61         |
| 2.7 Thought and Computer Exercises   | 63         |
| <b>3 Ant Colony Optimization Algorithms for the Traveling Salesman Problem</b> | <b>65</b>  |
| 3.1 The Traveling Salesman Problem   | 65         |
| 3.2 ACO Algorithms for the TSP   | 67         |
| 3.3 Ant System and Its Direct Successors                                       | 69         |
| 3.4 Extensions of Ant System   | 76         |
| 3.5 Parallel Implementations   | 82         |
| 3.6 Experimental Evaluation  | 84         |
| 3.7 ACO Plus Local Search  | 92         |
| 3.8 Implementing ACO Algorithms  | 99         |
| 3.9 Bibliographical Remarks  | 114        |
| 3.10 Things to Remember  | 117        |
| 3.11 Computer Exercises  | 117        |
| <b>4 Ant Colony Optimization Theory</b>  | <b>121</b> |
| 4.1 Theoretical Considerations on ACO  | 121        |
| 4.2 The Problem and the Algorithm  | 123        |
| 4.3 Convergence Proofs   | 127        |

|          |  |            |
|----------|--|------------|
| 4.4      | ACO and Model-Based Search   | 138        |
| 4.5      | Bibliographical Remarks  | 149        |
| 4.6      | Things to Remember   | 150        |
| 4.7      | Thought and Computer Exercises   | 151        |
| <b>5</b> | <b>Ant Colony Optimization for <math>\mathcal{NP}</math>-Hard Problems</b> | <b>153</b> |
| 5.1      | Routing Problems   | 153        |
| 5.2      | Assignment Problems  | 159        |
| 5.3      | Scheduling Problems  | 167        |
| 5.4      | Subset Problems  | 181        |
| 5.5      | Application of ACO to Other $\mathcal{NP}$ -Hard Problems                  | 190        |
| 5.6      | Machine Learning Problems  | 204        |
| 5.7      | Application Principles of ACO  | 211        |
| 5.8      | Bibliographical Remarks  | 219        |
| 5.9      | Things to Remember   | 220        |
| 5.10     | Computer Exercises   | 221        |
| <b>6</b> | <b>AntNet: An ACO Algorithm for Data Network Routing</b>                   | <b>223</b> |
| 6.1      | The Routing Problem  | 223        |
| 6.2      | The AntNet Algorithm   | 228        |
| 6.3      | The Experimental Settings  | 238        |
| 6.4      | Results  | 243        |
| 6.5      | AntNet and Stigmergy   | 252        |
| 6.6      | AntNet, Monte Carlo Simulation, and Reinforcement Learning                 | 254        |
| 6.7      | Bibliographical Remarks  | 257        |
| 6.8      | Things to Remember   | 258        |
| 6.9      | Computer Exercises   | 259        |
| <b>7</b> | <b>Conclusions and Prospects for the Future</b>                            | <b>261</b> |
| 7.1      | What Do We Know about ACO?   | 261        |
| 7.2      | Current Trends in ACO  | 263        |
| 7.3      | Ant Algorithms   | 271        |
|          | Appendix: Sources of Information about the ACO Field                       | 275        |
|          | References   | 277        |
|          | Index  | 301        |

# 1 From Real to Artificial Ants

*I am lost! Where is the line?!*

—*A Bug's Life*, Walt Disney, 1998

Ant colonies, and more generally social insect societies, are distributed systems that, in spite of the simplicity of their individuals, present a highly structured social organization. As a result of this organization, ant colonies can accomplish complex tasks that in some cases far exceed the individual capabilities of a single ant.

The field of “ant algorithms” studies models derived from the observation of real ants’ behavior, and uses these models as a source of inspiration for the design of novel algorithms for the solution of optimization and distributed control problems.

The main idea is that the self-organizing principles which allow the highly coordinated behavior of real ants can be exploited to coordinate populations of artificial agents that collaborate to solve computational problems. Several different aspects of the behavior of ant colonies have inspired different kinds of ant algorithms. Examples are foraging, division of labor, brood sorting, and cooperative transport. In all these examples, ants coordinate their activities via *stigmergy*, a form of indirect communication mediated by modifications of the environment. For example, a foraging ant deposits a chemical on the ground which increases the probability that other ants will follow the same path. Biologists have shown that many colony-level behaviors observed in social insects can be explained via rather simple models in which only stigmergic communication is present. In other words, biologists have shown that it is often sufficient to consider stigmergic, indirect communication to explain how social insects can achieve self-organization. The idea behind ant algorithms is then to use a form of *artificial stigmergy* to coordinate societies of artificial agents.

One of the most successful examples of ant algorithms is known as “ant colony optimization,” or ACO, and is the subject of this book. ACO is inspired by the foraging behavior of ant colonies, and targets discrete optimization problems. This introductory chapter describes how real ants have inspired the definition of artificial ants that can solve discrete optimization problems.

## 1.1 Ants’ Foraging Behavior and Optimization

The visual perceptive faculty of many ant species is only rudimentarily developed and there are ant species that are completely blind. In fact, an important insight of early research on ants’ behavior was that most of the communication among individuals, or between individuals and the environment, is based on the use of chemicals produced by the ants. These chemicals are called *pheromones*. This is different from,


for example, what happens in humans and in other higher species, whose most important senses are visual or acoustic. Particularly important for the social life of some ant species is the *trail pheromone*. Trail pheromone is a specific type of pheromone that some ant species, such as *Lasius niger* or the Argentine ant *Iridomyrmex humilis* (Goss, Aron, Deneubourg, & Pasteels, 1989), use for marking paths on the ground, for example, paths from food sources to the nest. By sensing pheromone trails foragers can follow the path to food discovered by other ants. This collective trail-laying and trail-following behavior whereby an ant is influenced by a chemical trail left by other ants is the inspiring source of ACO.

### 1.1.1 Double Bridge Experiments


The foraging behavior of many ant species, as, for example, *I. humilis* (Goss et al., 1989), *Linepithema humile*, and *Lasius niger* (Bonabeau et al., 1997), is based on indirect communication mediated by pheromones. While walking from food sources to the nest and vice versa, ants deposit pheromones on the ground, forming in this way a pheromone trail. Ants can smell the pheromone and they tend to choose, probabilistically, paths marked by strong pheromone concentrations.

The pheromone trail-laying and -following behavior of some ant species has been investigated in controlled experiments by several researchers. One particularly brilliant experiment was designed and run by Deneubourg and colleagues (Deneubourg, Aron, Goss, & Pasteels, 1990; Goss et al., 1989), who used a double bridge connecting a nest of ants of the Argentine ant species *I. humilis* and a food source. They ran experiments varying the ratio  $r = l_l/l_s$  between the length of the two branches of the double bridge, where  $l_l$  was the length of the longer branch and  $l_s$  the length of the shorter one.

In the first experiment the bridge had two branches of equal length ( $r = 1$ ; see figure 1.1a). At the start, 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 were observed over time. The outcome was that (see also figure 1.2a), although in the initial phase random choices occurred, eventually all the ants used the same 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 any of the branches. Yet, because of random fluctuations, a few more ants will select one branch over the other. Because ants deposit pheromone while walking, a larger number of ants on a branch results in a larger amount of pheromone on that branch; this larger amount of pheromone in turn stimulates more ants to choose that branch again, and so on until finally the ants converge to one

**Figure 1.1**

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


**Figure 1.2**

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

single path. This *autocatalytic* or *positive feedback* process is, in fact, an example of a self-organizing behavior of the ants: a macroscopic pattern (corresponding to the convergence toward one branch) emerges out of processes and interactions taking place at a “microscopic” level (Camazine, Deneubourg, Franks, Sneyd, Theraulaz, & Bonabeau, 2001; Haken, 1983; Nicolis & Prigogine, 1977). In our case the convergence of the ants’ paths to one branch represents the macroscopic collective behavior, which can be explained by the microscopic activity of the ants, that is, by the local interactions among the individuals of the colony. It is also an example of *stigmergic* communication (for a definition of *stigmergy*, see section 1.4): ants coordinate their activities, exploiting indirect communication mediated by modifications of the environment in which they move.

In the second experiment, the length ratio between the two branches was set to  $r = 2$  (Goss et al., 1989), so that the long branch was twice as long as the short one (figure 1.1b shows the experimental setup). In this case, in most of the trials, after some time all the ants chose to use only the short branch (see figure 1.2b). As in the first experiment, ants leave the nest to explore the environment and arrive at a decision point where they have to choose one of the two branches. Because the two branches initially appear identical to the ants, they choose randomly. 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, this experimental setup presents a remarkable difference with respect to the previous one: because one branch is shorter than the other (see figure 1.1b), the ants choosing the short branch are the first to reach the food and to start their return to the nest. But then, when they must make a decision between the short and the long branch, the higher level of pheromone on the short branch will bias their decision in its favor. Therefore, pheromone starts to accumulate faster on the short branch, which will eventually be used by all the ants because of the autocatalytic process described previously. When compared to the experiment with the two branches of equal length, the influence of initial random fluctuations is much reduced, and *stigmergy*, *autocatalysis*, and *differential path length* are the main mechanisms at work. Interestingly, it can be observed that, even when the long branch is twice as long as the short one, not all the ants use the short branch, but a small percentage may take the longer one. This may be interpreted as a type of “path exploration.”

It is also interesting to see what happens when the ant colony is offered, after convergence, a new shorter connection between the nest and the food. This situation was studied in an additional experiment in which initially only the long branch was

**Figure 1.3**

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

offered to the colony and after 30 minutes the short branch was added (see figure 1.3). In this case, the short branch was only selected sporadically and the colony was trapped on the long branch. This can be explained by the high pheromone concentration on the long branch and by the slow evaporation of pheromone. In fact, the great majority of ants choose the long branch because of its high pheromone concentration, and this autocatalytic behavior continues to reinforce the long branch, even if a shorter one appears. Pheromone evaporation, which could favor exploration of new paths, is too slow: the lifetime of the pheromone is comparable to the duration of a trial (Goss et al., 1989), which means that the pheromone evaporates too slowly to allow the ant colony to “forget” the suboptimal path to which they converged so that the new and shorter one can be discovered and “learned.”

### 1.1.2 A Stochastic Model

Deneubourg and colleagues (Deneubourg et al., 1990; Goss et al., 1989) proposed a simple stochastic model that adequately describes the dynamics of the ant colony as observed in the double bridge experiment. In this model,  $\psi$  ants per second cross the bridge in each direction at a constant speed of  $v$  cm/s, depositing one unit of pheromone on the branch. Given the lengths  $l_s$  and  $l_l$  (in cm) of the short and of the long branch, an ant choosing the short branch will traverse it in  $t_s = l_s/v$  seconds, while an ant choosing the long branch will use  $r \cdot t_s$  seconds, where  $r = l_l/l_s$ .

The probability  $p_{ia}(t)$  that an ant arriving at decision point  $i \in \{1, 2\}$  (see figure 1.1b) selects branch  $a \in \{s, l\}$ , where  $s$  and  $l$  denote the short and long branch respectively, at instant  $t$  is set to be a function of the total amount of pheromone  $\varphi_{ia}(t)$



on the branch, which is proportional to the number of ants that used the branch until time  $t$ . For example, the probability  $p_{is}(t)$  of choosing the short branch is given by

$$p_{is}(t) = \frac{(t_s + \varphi_{is}(t))^\alpha}{(t_s + \varphi_{is}(t))^\alpha + (t_s + \varphi_{il}(t))^\alpha}, \quad (1.1)$$

where the functional form of equation (1.1), as well as the value  $\alpha = 2$ , was derived from experiments on trail-following (Deneubourg et al., 1990);  $p_{il}(t)$  is computed similarly, with  $p_{is}(t) + p_{il}(t) = 1$ .

This model assumes that the amount of pheromone on a branch is proportional to the number of ants that used the branch in the past. In other words, no pheromone evaporation is considered by the model (this is in accordance with the experimental observation that the time necessary for the ants to converge to the shortest path has the same order of magnitude as the mean lifetime of the pheromone (Goss et al., 1989; Beckers, Deneubourg, & Goss, 1993)). The differential equations that describe the evolution of the stochastic system are

$$d\varphi_{is}/dt = \psi p_{js}(t - t_s) + \psi p_{is}(t), \quad (i = 1, j = 2; i = 2, j = 1), \quad (1.2)$$

$$d\varphi_{il}/dt = \psi p_{jl}(t - r \cdot t_s) + \psi p_{il}(t), \quad (i = 1, j = 2; i = 2, j = 1). \quad (1.3)$$

Equation (1.2) can be read as follows: the instantaneous variation, at time  $t$ , of pheromone on branch  $s$  and at decision point  $i$  is given by the ants' flow  $\psi$ , assumed constant, multiplied by the probability of choosing the short branch at decision point  $j$  at time  $t - t_s$  plus the ants' flow multiplied by the probability of choosing the short branch at decision point  $i$  at time  $t$ . The constant  $t_s$  represents a time delay, that is, the time necessary for the ants to traverse the short branch. Equation (1.3) expresses the same for the long branch, except that in this case the time delay is given by  $r \cdot t_s$ .

The dynamic system defined by these equations was simulated using the Monte Carlo method (Liu, 2001). In figure 1.4, we show the results of two experiments consisting of 1000 simulations each and in which the branch length ratio was set to  $r = 1$  and to  $r = 2$ . It can be observed that when the two branches have the same length ( $r = 1$ ) the ants converge toward the use of one or the other of the branches with equal probability over the 1000 simulations. Conversely, when one branch is twice as long as the other ( $r = 2$ ), then in the great majority of experiments most of the ants choose the short branch (Goss et al., 1989).

In this model the ants deposit pheromone both on their forward and their backward paths. It turns out that this is a necessary behavior to obtain convergence of the ant colony toward the shortest branch. In fact, if we consider a model in which ants deposit pheromone only during the forward or only during the backward trip, then