Université Libre de Bruxelles
Faculté des Sciences Appliquées

Ant Colony Optimization and its Application to Adaptive Routing in Telecommunication Networks

Gianni Di Caro

Dissertation présentée en vue de l'obtention du grade de
Docteur en Sciences Appliquées

Bruxelles, September 2004

Promoteur: Prof. Marco Dorigo
IRIDIA, Institut de Recherches Interdisciplinaires
et de Développements en Intelligence Artificielle

Année académique 2003-2004

Abstract

In *ant societies*, and, more in general, in insect societies, the activities of the individuals, as well as of the society as a whole, are not regulated by any explicit form of centralized control. On the other hand, adaptive and robust behaviors transcending the behavioral repertoire of the single individual can be easily observed at society level. These complex global behaviors are the result of self-organizing dynamics driven by local interactions and communications among a number of relatively simple individuals. The simultaneous presence of these and other fascinating and unique characteristics have made ant societies an attractive and inspiring model for building new *algorithms* and new *multi-agent systems*. In the last decade, ant societies have been taken as a reference for an ever growing body of scientific work, mostly in the fields of robotics, operations research, and telecommunications.

Among the different works inspired by ant colonies, the *Ant Colony Optimization metaheuristic* (ACO) is probably the most successful and popular one. The ACO metaheuristic is a multi-agent framework for combinatorial optimization whose main components are: a set of *ant-like agents*, the use of *memory* and of *stochastic decisions*, and strategies of *collective* and *distributed learning*. It finds its roots in the experimental observation of a specific foraging behavior of some ant colonies that, under appropriate conditions, are able to select the *shortest path* among few possible paths connecting their nest to a food site. The *pheromone*, a volatile chemical substance laid on the ground by the ants while walking and affecting in turn their moving decisions according to its local intensity, is the mediator of this behavior. All the elements playing an essential role in the ant colony foraging behavior were understood, thoroughly reverse-engineered and put to work to solve problems of combinatorial optimization by Marco Dorigo and his co-workers at the beginning of the 1990's. From that moment on it has been a flourishing of new combinatorial optimization algorithms designed after the first algorithms of Dorigo's *et al.*, and of related scientific events. In 1999 the ACO metaheuristic was defined by Dorigo, Di Caro and Gambardella with the purpose of providing a common framework for describing and analyzing all these algorithms inspired by the same ant colony behavior and by the same common process of reverse-engineering of this behavior. Therefore, the ACO metaheuristic was defined *a posteriori*, as the result of a synthesis effort effectuated on the study of the characteristics of all these ant-inspired algorithms and on the abstraction of their common traits. The ACO's synthesis was also motivated by the usually good performance shown by the algorithms (e.g., for several important combinatorial problems like the quadratic assignment, vehicle routing and job shop scheduling, ACO implementations have outperformed state-of-the-art algorithms).

The definition and study of the ACO metaheuristic is one of the two fundamental goals of the thesis. The other one, strictly related to this former one, consists in the design, implementation, and testing of ACO instances for problems of *adaptive routing in telecommunication networks*.

This thesis is an in-depth journey through the ACO metaheuristic, during which we have (re)defined ACO and tried to get a clear understanding of its potentialities, limits, and relationships with other frameworks and with its biological background. The thesis takes into account all the developments that have followed the original 1999's definition, and provides a formal and comprehensive systematization of the subject, as well as an up-to-date and quite comprehensive review of current applications. We have also identified in dynamic problems in telecommuni-

cation networks the most appropriate domain of application for the ACO ideas. According to this understanding, in the most applicative part of the thesis we have focused on problems of adaptive routing in networks and we have developed and tested four new algorithms.

Adopting an original point of view with respect to the way ACO was firstly defined (but maintaining full conceptual and terminological consistency), ACO is here defined and mainly discussed in the terms of *sequential decision processes* and *Monte Carlo sampling* and *learning*. More precisely, ACO is characterized as a policy search strategy aimed at learning the distributed parameters (called *pheromone variables* in accordance with the biological metaphor) of the stochastic decision policy which is used by so-called *ant* agents to generate solutions. Each ant represents in practice an independent *sequential decision process* aimed at *constructing* a possibly feasible solution for the optimization problem at hand by using only information *local* to the decision step. Ants are *repeatedly* and *concurrently* generated in order to sample the solution set according to the current policy. The outcomes of the generated solutions are used to *partially evaluate* the current policy, *spot* the most promising search areas, and *update the policy parameters* in order to possibly focus the search in those promising areas while keeping a satisfactory level of overall *exploration*.

This way of looking at ACO has facilitated to disclose the strict relationships between ACO and other well-known frameworks, like *dynamic programming*, *Markov* and *non-Markov decision processes*, and *reinforcement learning*. In turn, this has favored reasoning on the general properties of ACO in terms of amount of complete *state information* which is used by the ACO's ants to take optimized decisions and to encode in pheromone variables memory of both the decisions that belonged to the sampled solutions and their quality.

The ACO's biological context of inspiration is fully acknowledged in the thesis. We report with extensive discussions on the shortest path behaviors of ant colonies and on the identification and analysis of the few nonlinear dynamics that are at the very core of self-organized behaviors in both the ants and other societal organizations. We discuss these dynamics in the general framework of *stigmergic modeling*, based on asynchronous environment-mediated communication protocols, and (pheromone) variables priming coordinated responses of a number of "cheap" and concurrent agents.

The second half of the thesis is devoted to the study of the application of ACO to problems of *online routing in telecommunication networks*. This class of problems has been identified in the thesis as the most appropriate for the application of the multi-agent, distributed, and adaptive nature of the ACO architecture. Four novel ACO algorithms for problems of adaptive routing in telecommunication networks are throughly described. The four algorithms cover a wide spectrum of possible types of network: two of them deliver *best-effort traffic in wired IP networks*, one is intended for *quality-of-service (QoS) traffic in ATM networks*, and the fourth is for *best-effort traffic in mobile ad hoc networks*. The two algorithms for wired IP networks have been extensively tested by simulation studies and compared to state-of-the-art algorithms for a wide set of reference scenarios. The algorithm for mobile ad hoc networks is still under development, but quite extensive results and comparisons with a popular state-of-the-art algorithm are reported. No results are reported for the algorithm for QoS, which has not been fully tested. The observed experimental performance is excellent, especially for the case of wired IP networks: our algorithms always perform comparably or much better than the state-of-the-art competitors. In the thesis we try to understand the rationale behind the brilliant performance obtained and the good level of popularity reached by our algorithms. More in general, we discuss the reasons of the general efficacy of the ACO approach for network routing problems compared to the characteristics of more classical approaches. Moving further, we also informally define *Ant Colony Routing* (ACR), a multi-agent framework explicitly integrating learning components into the ACO's design in order to define a general and in a sense futuristic architecture for autonomic network control.

Most of the material of the thesis comes from a re-elaboration of material co-authored and published in a number of books, journal papers, conference proceedings, and technical reports. The detailed list of references is provided in the Introduction.

Contents