CPSC 609 Topic Paper - John Brosz

1 Ant Colony Optimization

Ant colony optimization (ACO) is a optimization technique that is good at finding a optimal solutions to difficult, discrete optimization problems [5]. These techniques are inspired by ant colony foraging behavior.

Research has shown that although ants are blind, they are able to find shortest path between food sources and the ant colony [2]. Ants accomplish this feat by laying down trails of pheromone. There are two key features about these trails. The first is that other ants, upon encountering these trails are probabilistically more likely to follow a trail (especially one with very concentrated pheromones) than to continue in their previous, essentially random, course. The second key feature is that the pheromones comprising trails evaporate as time passes [6]. To understand how this works, consider the scenario shown in Figure 1. If a number of ants, say fifty, start at the food and head towards the nest, if no pheromones trails exist, probabilistically equal numbers of ants should follow each path. However, the ants on the shortest path will reach the colony first, drop off their food, and head back for more. On their way back they will again have to choose between the two paths. However, at this point since more ants have completed the trip along the short path, this path possess a greater concentration of pheromones, making it more likely for returning ants to follow shorter path. Eventually most, if not all, ants will end up following the shorter path [1].

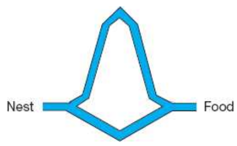

Figure 1: Two possible paths for ants from the nest to a food source. Image adapted from [1].

In ACO solutions are found by using a group of independant, cooperating agents (artificial ants) to explore solutions to a given optimization problem [5]. The artificial ants (AAs) have several similarities to real real ants. The first is that AAs cooperate as a colony to find optimal solutions. Another similarity is that AAs use the concept of stigmergy, akin to the dropping of pheromones by ants, to communicate and find the best solution. The artificial pheromone trails left by the AAs also disperse over time like real ant pheromone trails [4]. Lastly AAs use a myopic and stochastic transition decision policy to move between adjacent locations (or states) within the problem environment [5].

There are also several differences between real ants and AAs. AAs have an internal state that can be used to store past actions to influence current decisions [6]. Artificial ants may deposit pheromone while they are constructing a solution (e.g., a path through the problem) or after they have completed an entire solution [1]. Furthermore, the amount of pheromone deposited along the path can be a function of the solution's quality [6]. Lastly most ACOs give AAs extra knowledge such as local optimization information to assist in improving the solution [5].

The first ACO system and many subsequent variations were developed to solve the traveling salesman problem [4]. Some of the other problems ACOs have been used to solve are the asymmetric traveling salesman problem, sequential ordering problem, quadratic assignment problem, several vehicle routing problems, job and bus scheduling problems, graph coloring problem, graph partitioning problem, and with finding optimal routing paths in telecommunications networks [4].

ACO is quite interesting as it possesses several unique properties that make these techniques optimal and robust. The first is that pheromone trail building is an autocatalytic process. The positive feedback loop can very quickly yield optimal results although stagnation (convergence to non-optimal solutions) be must avoided [5]. Generally ACOs avoid stagnation through pheromone evaporation and with the probabilistic behavior of the ants. ACOs also are able effectively to handle dynamic problems, such as dealing with routing paths, where paths or nodes of the graph may be suddenly changed or removed. This is due to the manner in which ACOs ensure that even non-optimal paths are occasionally traversed by ants[1]. The use of stigmergy and independant agents make ACO techniques good candidates for parallelization [6]. In terms of performance ACOs perform competitively for the problems previously mentioned but are sometimes outperformed by problem-specific algorithms [10].

2 Reference Justification

One guide for selection of articles that I have used, is the number of times each article has been listed as being cited by http://scholar.google.com. These numbers are presented in brackets at the end of each work's justification.

The Ant Colony Optimization Website [4] maintained by Marco Dorigo is an important reference because is provides an overview of ACO as well as gives a guide as to which ACO papers may be useful.

Dorigo, Maniezzo, and Colorni's work [6] is listed because it presents the first ACO, the Ant System. Dorigo's PhD thesis, an earlier publication on the Ant System, is not listed here because it is only available in Italian. (1073).

Dorigo and Gambardella's work [3] is interested because it provides in depth testing and analysis of many of the variables and constants used within the Ant System. (179).

Dorigo, Di Caro, and Gambardella's work [5] provides another relatively recent overview of existing ant colony optimization techniques. It also presents ACO as a problem solving heuristic. (385).

The *Nature* article [1] provides a good, concise, and relatively recent overview of ant colony optimization systems. (145).

The ant-based control system [9] provides another ant optimization system that, as far as I can tell, was derived independently of Dorigo et al's work. This system was designed for dynamic telecommunications routing (rather than the traveling salesman problem). (250).

The work of Stutzle and Hoos [10] introduces a variation on the Ant System that improves ant system performance to levels comparable with other fine-tuned traveling salesman and quadratic assignment problem algorithms. (130).

Gambardella, Taillard, and Agazzi [8] have created an interesting ant colony optimization system that makes use of two ant colonies to solve the vehicle routing problem with time windows. This system performs competitively with other known methods [8]. (124).

Deneubourg and Goss's contribution [2] is the most frequently referenced study of actual ant behavior I encountered in the ant optimization literature. (166).

Dorigo and Stutzle's book [7] is included because it is likely the most comprehensive and thorough overview of ACOs and their applications. (93).

References

[1] E. Bonabeau, M. Dorigo, and G. Theraulaz. Inspiration for optimization from social insect behaviour. *Nature*, 406, July 2000.

[2] J. L. Deneubourg and S. Goss. Collective patterns and decision making. *Ethology, Ecology, and Evolution*, 1:295–311, 1989.

[3] M. Dorigo and L. M. Gambardella. Ant colonies for the traveling salesman problem. *BioSystems*, 43:73–81, 1997.

[4] Marco Dorigo. Ant colony optimization. URL http://www.aco-metaheuristic.org, 2004. Accessed Jan 2006.

[5] Marco Dorigo, Gianni Di Caro, and Luca M. Gambardella. Ant algorithms for discrete optimizations. *Artificial Life*, 5(2):137–172, 1999.

[6] Marco Dorigo and Alberto Colorni. The ant system: Optimization by a colony of cooperating agents. *IEEE Transactions on Systems, Man, and Cybernetics - Part B*, 26(1):1–13, 1996.

[7] Marco Dorigo and Thomas Stutzle. *Ant Colony Optimization*. MIT Press, June 2004.

[8] Luca Maria Gambardella, Eric Taillard, and Giovanni Agazzi. Macs-vrptw: a multiple ant colony system for vehicle routing problems with time windows. *New ideas in optimization*, pages 63–76, 1999.

[9] Ruud Schoonderwoerd, Janet L. Bruten, Owen E. Holland, and Leon J. M. Rothkrantz. Ant-based load balancing in telecommunications networks. *Adaptive Behavior*, 5(2):169–207, 1996.

[10] T. Stutzle and H. Hoos. Max-min ant system. *Future Generation Computer Systems*, 16(8):889–914, 2000.