

An Ant Colony Optimization Approach to the Probabilistic Traveling Salesman Problem

Leonora Bianchi¹, Luca Maria Gambardella¹ and Marco Dorigo²

¹ IDSIA, Strada Cantonale Galleria 2, CH-6928 Manno, Switzerland
{leonora,luca}@idsia.ch
<http://www.idsia.ch>

² Université Libre de Bruxelles, IRIDIA,
Avenue Franklin Roosevelt 50, CP 194/6, 1050 Bruxelles, Belgium
mdorigo@ulb.ac.be

Abstract. The Probabilistic Traveling Salesman Problem (PTSP) is a TSP problem where each customer has a given probability of requiring a visit. The goal is to find an a-priori tour of minimal expected length over all customers, with the strategy of visiting a random subset of customers in the same order as they appear in the a-priori tour.

We address the question of whether and in which context an a-priori tour found by a TSP heuristic can also be a good solution for the PTSP. We answer this question by testing the relative performance of two ant colony optimization algorithms, Ant Colony System (ACS) introduced by Dorigo and Gambardella for the TSP, and a variant of it (pACS) which aims to minimize the PTSP objective function.

We show in which probability configuration of customers pACS and ACS are promising algorithms for the PTSP.

1 Introduction

Consider a routing problem through a set V of n customers. On any given instance of the problem each customer i has a known position and a probability p_i of actually requiring a visit, independently of the other customers. Finding a solution for this problem implies having a strategy to determine a tour for each random subset $S \subseteq V$, in such a way as to minimize the expected tour length. The most studied strategy is the a-priori one. An a-priori strategy has two components: the a-priori tour and the updating method. The a-priori tour is a tour visiting the complete set V of n customers; the updating method modifies the a-priori tour in order to have a particular tour for each subset of customers $S \subseteq V$. A very simple example of updating method is the following: for every subset of customers, visit them *in the same order* as they appear in the a-priori tour, skipping the customers that do not belong to the subset. The strategy related to this method is called the 'skipping strategy'. The problem of finding an a-priori tour of minimum expected length under the skipping strategy is defined as the Probabilistic Traveling Salesman Problem (PTSP). This is an NP-hard problem [3, 1], and was introduced in Jaillet's PhD thesis [8].

The PTSP approach models applications in a delivery context where a set of customers has to be visited on a regular (e.g., daily) basis, but all customers do not always require a visit, and where re-optimizing vehicle routes from scratch every day is unfeasible. In this context the delivery man would follow a standard route (i.e., an a-priori tour), leaving out customers that on that day do not require a visit. The standard route of least expected length corresponds to the optimal PTSP solution.

In the literature there are a number of algorithmic and heuristic approaches used to find suboptimal solutions for the PTSP. Heuristics using a nearest neighbor criterion or savings criterion were implemented and tested by Jézéquel [9] and by Rossi-Gavioli [11]. Later, Bertsimas-Jaillet-Odoni [3] and Bertsimas-Howell [2] have further investigated some of the properties of the PTSP and have proposed some heuristics for the PTSP. These include tour construction heuristics (space filling curves and radial sort), and tour improvement heuristics (probabilistic 2-opt edge interchange local search and probabilistic 1-shift local search). Most of the heuristics proposed are an adaptation of a TSP heuristics to the PTSP, or even the TSP heuristics itself, which in some cases gives good PTSP solutions. More recently, Laporte-Louveaux-Mercure [10] have applied an integer L-shaped method to the PTSP and have solved to optimality instances involving up to 50 vertices.

No application of nature-inspired algorithms such as ant colony optimization (ACO) [4] or genetic algorithms has been done yet. This paper investigates the potentialities of ACO algorithms for the PTSP. In the remainder of the paper we first introduce the PTSP objective function and notations (section 2), then we describe the ACO algorithms which we tested (section 3), and we finally show the experimental results obtained (section 4).

2 The PTSP objective function

Let us consider an instance of the PTSP. We have a complete connected graph whose nodes form a set $V = \{i = 1, 2, \dots, n\}$ of customers. Each customer has a probability p_i of requiring a visit, independently from the others. A solution for this instance is a tour τ over all nodes in V (an ‘a-priori tour’), to which is associated the expected length objective function

$$E[L_\tau] = \sum_{S \subseteq V} p(S) L_\tau(S) . \quad (1)$$

In the above expression S is a subset of the set of nodes V , $L_\tau(S)$ is the distance required to visit the subset of customers S (in the same order as they appear in the a-priori tour), and $p(S)$ is the probability for the subset of customers S to require a visit:

$$p(S) = \prod_{i \in S} p_i \prod_{i \notin S} (1 - p_i) . \quad (2)$$

Jaillet [8] showed that the evaluation of the PTSP objective function (eq.(1)) can be done in $O(n^2)$. In fact, let us consider (without loss of generality) an

a-priori tour $\tau = (1, 2, \dots, n)$; then its expected length is

$$E[L_\tau] = \sum_{i=1}^n \sum_{j=i+1}^n d_{ij} p_i p_j \prod_{k=i+1}^{j-1} (1 - p_k) + \sum_{i=1}^n \sum_{j=1}^{i-1} d_{ij} p_i p_j \prod_{k=i+1}^n (1 - p_k) \prod_{l=1}^{j-1} (1 - p_l) . \quad (3)$$

This expression is derived by looking at the probability for each couple of customers to require a visit. In the special class of PTSP instances where $p_i = p$ for all customers $i \in V$ (the homogeneous PTSP), equation (3) becomes

$$E[L_\tau] = p^2 \sum_{r=0}^{n-2} (1 - p)^r L_\tau^{(r)} \quad (4)$$

where $L_\tau^{(r)} \equiv \sum_{j=1}^n d(j, (j + 1 + r) \bmod n)$. The $L_\tau^{(r)}$'s have the combinatorial interpretation of being the lengths of $\gcd(n, r + 1)$ sub-tours¹, obtained from tour τ by visiting one point and skipping the next r points.

3 Ant Colony Optimization

In ACO algorithms a colony of artificial ants iteratively constructs solutions for the problem under consideration using artificial pheromone trails and heuristic information. The pheromone trails are modified by ants during the algorithm execution in order to store information about ‘good’ solutions. Most ACO algorithms follow the algorithmic scheme given in Figure 1.

```

procedure ACO metaheuristic for combinatorial optimization problems
  Set parameters, initialize pheromone trails
  while (termination condition not met)
    ConstructSolutions
    (ApplyLocalSearch)
    UpdateTrails
  end while

```

Fig. 1. High level pseudocode for the ACO metaheuristic

ACO are solution construction algorithms, which, in contrast to local search algorithms, may not find a locally optimal solution. Many of the best performing ACO algorithms improve their solutions by applying a local search algorithm after the solution construction phase. Our primary goal in this work is to analyze the PTSP tour construction capabilities of ACO, hence in this first investigation we do not use local search.

¹ The term ‘gcd’ stays for ‘greatest common divisor’.

We apply to the PTSP Ant Colony System (ACS) [7, 5], a particular ACO algorithm which was successfully applied to the TSP. We also consider a modification of ACS which explicitly takes into account the PTSP objective function (we call this algorithm probabilistic ACS, that is, pACS). In the following, we describe how ACS and pACS build a solution and how they update pheromone trails.

3.1 Solution construction in ACS and pACS

A feasible solution for an n -city PTSP is an a-priori tour which visits all customers. Initially m ants are positioned on their starting cities chosen according to some initialization rule (e.g., randomly). Then, when an ant k is on city i , the next city is chosen as follows.

- With probability q_0 , a city j that maximizes $\tau_{ij} \cdot \eta_{ij}^\beta$ is chosen in the set $J_k(i)$ of the cities not yet visited by ant k . Here, τ_{ij} is the pheromone trail associated to the arc joining city i and city j , η_{ij} is the inverse distance $1/d_{ij}$ between city i and city j , and β is a parameter which determines the influence of the heuristic information.
- With probability $1 - q_0$, city j is chosen randomly with a probability given by

$$p_k(i, j) = \begin{cases} \frac{\tau_{ij} \cdot \eta_{ij}^\beta}{\sum_{r \in J_k(i)} \tau_{ir} \cdot \eta_{ir}^\beta}, & \text{if } j \in J_k(i) \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

Hence, with a probability q_0 the ant chooses the best city according to the pheromone trail and to the distance between cities, while with probability $1 - q_0$ it explores the search space in a biased way.

3.2 Pheromone trails update in ACS and pACS

Pheromone trails are updated in two stages. The first stage of pheromone update is done step-by-step during the tour construction phase after each ant has chosen the next city. This is done according to the local updating rule $\tau_{ij} \leftarrow (1 - \rho) \cdot \tau_{ij} + \rho \cdot \tau_0$, where ρ , $0 < \rho \leq 1$, and τ_0 , are two parameters. The effect of the local updating rule is to make less desirable to other ants an arc which has already been chosen by an ant, so that the exploration of different tours is favored during one iteration of the algorithm.

The second stage of pheromone update occurs when all ants have terminated their tour. Pheromone is modified on those arcs belonging to the best tour since the beginning of the trial (global-best tour), by the following global updating rule

$$\tau_{ij} \leftarrow (1 - \alpha) \cdot \tau_{ij} + \alpha \cdot \Delta\tau_{ij}, \quad (6)$$

where

$$\Delta\tau_{ij} = \text{ObjectiveFunc}_{best}^{-1} \quad (7)$$

with $0 < \alpha \leq 1$ being the pheromone decay parameter, and $ObjectiveFunc_{best}$ is the value of the objective function of the global-best tour. In ACS the objective function is the a-priori tour length, while in pACS the objective function is the PTSP expected length of the a-priori tour. In the next section we discuss in more detail the differences between ACS and pACS.

3.3 Discussion of differences between ACS and pACS

Differences between ACS and pACS are due to the fact that the two algorithms exploit different objective functions in the pheromone updating phase. The global updating rule of equations (6) and (7) implies two differences in the way ACS and pACS explore the search space. The first and most important difference is the set of arcs on which pheromone is globally increased, which is in general different in ACS and pACS. In fact, the ‘best tour’ in eq. (7) is relative to the objective function. Therefore in ACS the search will be biased toward the shortest tour, while in pACS it will be biased toward the least expected length tour. The second difference between ACS and pACS is in the quantity $\Delta\tau_{ij}$ by which pheromone is increased on the selected arcs. This aspect is less important than the first, because ACO in general are more sensitive to the difference of pheromone among arcs than to its absolute value.

The length of an a-priori tour (ACS objective function) may be considered as an $O(n)$ approximation to the $O(n^2)$ expected length (pACS objective function). In general, the worse the approximation, the worse will be the solution quality of ACS versus pACS. The quality of the approximation depends on the set of customer probabilities p_i . In the homogeneous PTSP, where customer probability is p for all customers, it is easy to see the relation between the two objective functions. For a given a-priori tour L_τ we have

$$\Delta = L_\tau - E[L_\tau] = (1 - p^2)L_\tau - \sum_{r=1}^{n-2} (1 - p)^r L_\tau^{(r)}, \quad (8)$$

which implies

$$\Delta \sim O(q \cdot L_\tau) \quad (9)$$

for $(1 - p) = q \rightarrow 0$. Therefore the higher the probability, the better is the a-priori tour length L_τ as an approximation for the expected length $E[L_\tau]$.

In the heterogeneous PTSP, it is not easy to see the relation between the two objective functions, since each arc of the a-priori tour L_τ is multiplied by a different probabilistic weight (see eq.(3)), and a term with L_τ cannot be isolated in the expression of $E[L_\tau]$, as in the homogeneous case.

ACS and pACS also differ in time complexity. In both algorithms one iteration (i.e., one cycle through the *while* condition of Fig. 1 is $O(n^2)$ [6], but the constant of proportionality is bigger in pACS than in ACS. To see this one should consider the procedure *UpdateTrail* of Fig. 1, where the global-best tour must be evaluated in order to choose the arcs on which pheromone is to be updated.

The global-best tour evaluation requires an $O(n)$ and $O(n^2)$ objective function calculation respectively in ACS and pACS. ACS is thus faster and always performs more iterations than pACS for a fixed CPU time.

4 Experimental tests

4.1 Homogeneous PTSP Instances

Homogeneous PTSP instances were generated starting from TSP instances and assigning to each customer a probability p of requiring a visit. TSP instances were taken from two benchmarks. The first is the TSPLIB at <http://www.iwr.uni-heidelberg.de/groups/comopt/software/TSPLIB95/tsp/>. From this benchmark we considered instances with a number of city between 50 and 200. The second benchmark is a group of instances where customers are randomly distributed on the square $[0, 10^6]$. Both uniform and clustered distributions were considered in this case, and the number of cities varied between 50 and 350. For generating random instances we used the Instance Generator Code of the 8th DIMACS Implementation Challenge at <http://research.att.com/dsj/chtsp/download.html>.

4.2 Computational environment and ACS parameters

Experiments were run on a Pentium Xeon, 1GB of RAM, 1.7 GHz processor. In order to assess the relative performance of ACS versus pACS independently from the details of the settings, the two algorithms were run with the same parameters. We chose the same settings which yielded good performance in earlier studies with ACS on the TSP [5]: 10 ants, $\beta = 2$, $q_0 = 0.98$, $\alpha = \rho = 0.1$ and $\tau_0 = 1/(n \cdot Obj)$, where n is the number of customers and Obj is the value of the objective function evaluated with the nearest neighbor heuristic [5].

The stopping criterion used in both algorithms is CPU time in seconds, chosen according to the relation $stoptime = k \cdot n^2$, with $k = 0.01$. This value of k lets ACS perform at least $17 \cdot 10^3$ iterations on problems up to 100 customers, and at least $15 \cdot 10^3$ iterations on problems with more than 100 customers.

4.3 Results

For each TSP instance tested, nine experiments were done varying the value of the customer probability p from 0.1 to 0.9 with a 0.1 interval. Fig. 2 summarizes results obtained on 21 symmetric PTSP instances, one third of the instances where taken from the TSPLIB, the others were random instances (half of them uniform and half clustered). The figure shows the relative performance of pACS versus ACS, averaged over the tested instances. A typical result for a single instance is reported in Fig. 3.

As it was reasonable to expect, for small enough probabilities pACS outperforms ACS. In all problems we tested though, there is a range of probabilities $[p_0, 1]$ for which ACS outperforms pACS. The critical probability p_0 at which this happens depends on the problem.




Fig. 2. Relative performance of pACS versus ACS for the homogeneous PTSP. The vertical axis represents $(E[L_\tau(pACS)] - E[L_\tau(ACS)]) / E[L_\tau(pACS)]$. On the horizontal axis there is the customer probability p . Each point of the graph is an average over 21 symmetric homogeneous PTSP instances. Error bars represent average deviation, defined as $\sum_{i=1}^n |x_i - \langle x \rangle| / n$. Note that for $p = 1$ ACS outperforms pACS, since for a fixed CPU stopping time ACS makes more iterations




Fig. 3. Relative performance of pACS versus ACS for the eil76.tsp instance of the TSPLIB. Error bars represent the average deviation of the result over 5 independent runs of the algorithms

The reason why pACS does not always perform better than ACS is clear if we consider two aspects: the complexity (speed) of ACS versus pACS, and the goodness of the PTSP objective function approximation as p approaches 1. When p is near to 1, a good solution to the TSP is also a good solution to the PTSP, therefore ACS, which performs more iterations than pACS, has a better chance to find a good solution.

4.4 Absolute performance

The optimal solution of the PTSP instances we tested is not known, therefore an absolute performance evaluation of the pACS heuristic can only be done against a theoretical lower bound, when this is available and tight enough. A lower bound to the optimal solution would give us an upper bound to the error performed by pACS heuristic. In fact, if LB is the lower bound and $E[L_{\tau^*}]$ is the optimal solution, then by definition we have

$$E[L_{\tau^*}] \geq LB . \quad (10)$$

If the solution value of pACS is $E[L_{\tau}]$, then the following inequality holds for the relative error

$$\frac{E[L_{\tau}] - E[L_{\tau^*}]}{E[L_{\tau^*}]} \leq \frac{E[L_{\tau}] - LB}{LB} . \quad (11)$$

For the homogeneous PTSP and for instances where the optimal length L_{TSP} of the corresponding TSP is known, it is possible to use the following lower bound to the optimal expected length [2]

$$LB = pL_{TSP}(1 - (1 - p)^{n-1}) . \quad (12)$$

If we put this lower bound into the right side of equation (11), we obtain an upper bound of the relative error of pACS. The relative error upper bound is significantly tight (less than 100%) only for probabilities bigger than 0.5. Fig. 4 shows the absolute performance of pACS, evaluated with this method, for a few TSPLIB instances.

5 Conclusions and future work

In this paper we investigate the potentialities of ACO algorithms for the PTSP. In particular we show that the pACS algorithm is a promising heuristic for homogeneous TSP instances. Moreover, for near one probabilities the TSP heuristic ACS is a better alternative to pACS.

At present we are investigating the heterogeneous PTSP, for different probability configurations of customers. This is an interesting direction of research, since it is nearer to a real-world problem than the homogeneous PTSP. We are also trying to improve pACS performance by inserting in the ants' tour construction criterion information about the customers probabilities. Work which will follow this paper also comprehends a comparison of pACS with respect to




Fig. 4. Upper bound of relative percent error of pACS for 5 TSPLIB instances. The horizontal axis represents the customer probability

other PTSP algorithms. Moreover, pACS should be improved by adding to the tour construction phase a local search algorithm. The best choice and design of such a local search is also an interesting issue for the PTSP.

References

1. D. J. Bertsimas. *Probabilistic Combinatorial Optimization Problems*. PhD thesis, MIT, Cambridge, MA, 1988.
2. D. J. Bertsimas and L. Howell. Further results on the probabilistic traveling salesman problem. *European Journal of Operational Research*, 65:68–95, 1993.
3. D. J. Bertsimas, P. Jaillet, and A. Odoni. A priori optimization. *Operations Research*, 38:1019–1033, 1990.
4. M. Dorigo, G. Di Caro, and L. M. Gambardella. Ant algorithms for discrete optimization. *Artificial Life*, 5(2):137–172, 1999.
5. M. Dorigo and L. M. Gambardella. Ant colony system: A cooperative learning approach to the traveling salesman problem. *IEEE Transactions on Evolutionary Computation*, 1(1):53–66, 1997.
6. M. Dorigo, V. Maniezzo, and A. Colomi. 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. L. Gambardella and M. Dorigo. Solving symmetric and asymmetric tsps by ant colonies. In *Proceedings of the IEEE Conference on Evolutionary Computation, ICEC96, Nagoya, Japan, May 20-22*, pages 622–627, 1996.
8. P. Jaillet. *Probabilistic Traveling Salesman Problems*. PhD thesis, MIT, Cambridge, MA, 1985.
9. A. Jézéquel. Probabilistic vehicle routing problems, 1985.

10. G. Laporte, F. Louveaux, and H. Mercure. An exact solution for the a priori optimization of the probabilistic traveling salesman problem. *Operations Research*, 42:543–549, 1994.
11. F. A. Rossi and I. Gavioli. *Aspects of Heuristic Methods in the Probabilistic Traveling Salesman Problem*, pages 214–227. World Scientific, Singapore, 1987.