Université Libre de Bruxelles

Institut de Recherches Interdisciplinaires
et de Développements en Intelligence Artificielle

An Experimental Study of Estimation-based Metaheuristics for the Probabilistic Traveling Salesman Problem

Prasanna BALAPRAKASH, Mauro BIRATTARI,
Thomas STÜTZLE, and Marco DORIGO

IRIDIA – Technical Report Series

Technical Report No.
TR/IRIDIA/2007-021

December 2007

IRIDIA – Technical Report Series
ISSN 1781-3794

Published by:
IRIDIA, *Institut de Recherches Interdisciplinaires et de Développements en Intelligence Artificielle*
Université Libre de Bruxelles
Av F. D. Roosevelt 50, CP 194/6
1050 Bruxelles, Belgium

Technical report number TR/IRIDIA/2007-021

An Experimental Study of Estimation-based Metaheuristics for the Probabilistic Traveling Salesman Problem

Prasanna Balaprakash pbalapra@ulb.ac.be
Mauro Birattari mbiro@ulb.ac.be
Thomas Stützle stuetzle@ulb.ac.be
Marco Dorigo mdorigo@ulb.ac.be

IRIDIA, CoDE, Université Libre de Bruxelles, Brussels, Belgium

December 17, 2007

Abstract

The probabilistic traveling salesman problem (PTSP), a paradigmatic example of a stochastic combinatorial optimization problem, is used to study routing problems under uncertainty. Recently, we introduced a new estimation-based iterative improvement algorithm for the PTSP and we showed that it outperforms for a number of instance classes the previous state-of-the-art algorithms. In this paper, we integrate this estimation-based iterative improvement algorithm into some metaheuristics to solve the PTSP and we study their performance.

1 Introduction

Optimization problems that are both combinatorial and stochastic are important in a large number of practically relevant contexts. In order to tackle these problems, it is customary to consider a setting in which the cost of each solution is a random variable, and the goal is to find a solution that minimizes some statistics of the latter. For a number of practical and theoretical reasons, the optimization is performed with respect to the expectation. The most widely used approach for this task is *empirical estimation*, where the expectation is estimated through Monte Carlo simulation [1].

The probabilistic traveling salesman problem (PTSP) [2] is a paradigmatic example of a stochastic combinatorial optimization problem. The PTSP is similar to the TSP with the difference that each node has a probability of requiring a visit. The *a priori* optimization approach for the PTSP consists in finding an *a priori* solution that visits all the nodes and that minimizes the expected cost of *a posteriori* solutions: The *a priori* solution must be found prior to knowing which nodes are to be visited. The associated *a posteriori* solution is computed *after* knowing which nodes need to be visited. It is obtained by skipping the nodes that do not require a visit and visiting the others in the order in which they appear in the *a priori* solution.

In this paper, we integrate our new state-of-the-art PTSP iterative improvement algorithm [3] into some metaheuristics that have been developed for solving the TSP. The key issue that we want to address is the following: given that the PTSP is similar to the TSP, do the results reported in the TSP literature transfer to the PTSP? Is the knowledge obtained from TSP optimization algorithms useful for solving the PTSP? A secondary motivation for our work comes from the fact that there is no conclusive evidence on which are the most appropriate metaheuristics for the PTSP. As pointed out by Bianchi in [4], the main reason for this situation is a lack of a rigorous experimental comparison among different metaheuristics—most of the papers focus on a single metaheuristic which is compared either to a simple heuristic or to exact methods when available.

For the experimental study, we have selected five most widely used metaheuristics such as random restart local search (RR-LS), simulated annealing (SA), iterated local search (ILS), ant colony optimization (ACO), and memetic algorithm (MA). We use the state-of-the-art PTSP iterative improvement algorithm as the underlying heuristic for all these algorithms and we compare them under two practically relevant settings.

2 Estimation-based metaheuristics

The adoption of a fast subsidiary local search is of crucial importance for an effective metaheuristic for the PTSP. We adopt `2.5-opt-EEais`, the state-of-the-art iterative improvement for the PTSP [3] as the underlying local search. `2.5-opt-EEais` is an estimation-based local search in which the cost differences among neighboring solutions are estimated through Monte Carlo evaluation. The effectiveness of this approach can be attributed to several variance reduction techniques such as the method of common random numbers, adaptive sample size, and importance sampling.

TSP-specific metaheuristics can be easily extended for solving the PTSP by using the estimation-based approach for computing the solution cost. This approach is widely known as sample average approximation [5]. In the estimation-based approach, since solution costs are estimated, there is uncertainty in conclusively determining if one solution is better than another. A naive strategy to handle this issue consists in using a large number of samples to estimate the cost of each solution with a low variance. This could mean that very few solutions will be explored due to the time required to estimate the cost of each solution. However, this issue can be alleviated by using a statistical test that allows the metaheuristic to select the most appropriate number of samples.

Metahuristics such as SA, RR-LS, and ILS adopt the same adaptive sample size procedure as in `2.5-opt-EEais`. Therefore, let us first focus on the usage of the adaptive sample size procedure in `2.5-opt-EEais`. At *each step* of the iterative improvement, *Student's t-test* is used to select the appropriate number of samples: Given two solutions, the cost difference is sequentially computed on a number of samples. As soon as the *t-test* rejects the null hypothesis that the value of the cost difference estimator is equal to zero, the computation is stopped. If no statistical evidence is gathered, then the computation is continued until a maximum number M of realizations is considered, where M is a parameter of the algorithm. The sign of the estimate determines the solution of lower cost. The implementation of SA is similar to `2.5-opt-EEais` except the fact that the former accepts worse solutions than the current solution according to a probabilistic acceptance criterion. In RR-LS and ILS, at *each iteration*, we used the same *Student's t-test* for comparing two solutions.

We have also investigated the adoption of a nonparametric statistical test, the *Wilcoxon test*, instead of *Student's t-test*. The results showed that although there is no significant difference in the solution quality, the adoption of *Wilcoxon test* needs more samples. This is attributed to the conservative nature of the nonparametric test.

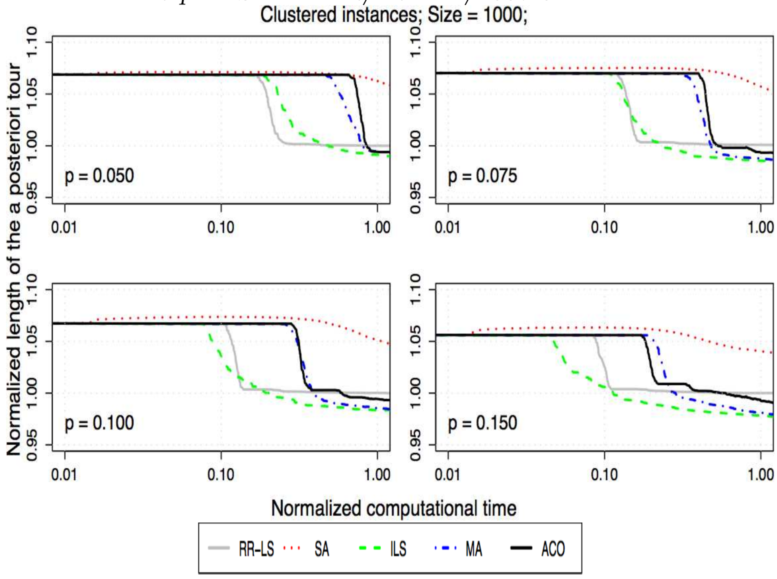

Figure 1: Experimental results on clustered homogeneous PTSP instances of size 1000. The plot represents the cost of the solutions obtained by SA, ILS, ACO, and MA normalized by the one obtained by random RR-LS. Each algorithm is run for 100 seconds.

In ant colony optimization and memetic algorithm, the selection-of-the-best is performed using a racing procedure called `F-race` [6]. This procedure is based on the sequential usage of the Friedman test. The adoption of the this procedure is motivated by the fact that, for the PTSP, previous studies have shown that this familywise nonparametric statistical test performs significantly better than the simple *Student's t-test* [7].

3 Preliminary experimental results

For the experimental analysis, we used homogeneous PTSP instances, where all the nodes of an instance have a same probability p of requiring a visit. The goal of our experiments is to assess the performance of the algorithms based on two factors: the variance of the cost estimator (the lower the value of p, the higher the variance) and the computation time. We considered ten levels for the former and three levels for the latter: short (100 seconds), medium (1000 seconds), and long (10000 seconds). All the parameters of the algorithms are fine tuned with a parameter tuning algorithm, `Iterative F-race` [8], for each level of the computation time. In this paper, we focus only on the performance of the algorithms observed under short computation times. In Figure 1, we report example results obtained on 100 *clustered homogeneous* PTSP instances of 1000 nodes for each of 4 levels of probability.

From the preliminary results, we can make the following observations: For the short computation time, the average cost of the solutions obtained by ILS and MA are better than all other algorithms. The average cost of the solutions obtained by ILS and MA are 1% lower than that of ACO and about 1% to 3% lower than the one of RR-LS. Note that these differences are significant in statistical sense according to a paired Wilcoxon test (α=0.05). The average solution cost of ILS is comparable to the one of MA. In spite of being provably convergent, SA is not practically useful: the average cost of the solutions obtained by this algorithm is approximately 5% higher than the one of RR-LS.

4 Conclusion and future work

In this paper, we integrated a new state-of-the-art PTSP local search into TSP specific metaheuristics for tackling the PTSP. The preliminary results show that the performance of the algorithms are quite similar to the one obtained in the context of TSP. This gives a clear indication that the knowledge obtained from solving the TSP can be transferred for tackling the PTSP. As in the TSP, the memetic algorithm and iterated local search emerge as high performing algorithms. Further research will be devoted to assess the performance of the proposed algorithms in a more detailed experimental setting.

References

[1] Fu, M.C.: Optimization via simulation: A review. Annals of Operations Research **53** (1994) 199–248

[2] Jaillet, P.: Probabilistic Traveling Salesman Problems. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA (1985)

[3] Balaprakash, P., Birattari, M., Stützle, T., Dorigo, M.: Adaptive sample size and importance sampling in estimation-based local search for stochastic combinatorial optimization: A complete analysis. Technical Report TR/IRIDIA/2007-015, IRIDIA, Université Libre de Bruxelles, Brussels, Belgium (2007)

[4] Bianchi, L.: Ant Colony Optimization and Local Search for the Probabilistic Traveling Salesman Problem: A Case Study in Stochastic Combinatorial Optimization. PhD thesis, Université Libre de Bruxelles, Brussels, Belgium (2006)

[5] Kleywegt, A., Shapiro, A., de Mello, T.H.: The sample average approximation method for stochastic discrete optimization. SIAM Journal on Optimization **12**(2) (2002) 479–502

[6] Birattari, M.: The Problem of Tuning Metaheuristics as Seen from a Machine Learning Perspective. PhD thesis, Université Libre de Bruxelles, Brussels, Belgium (2004)

[7] Birattari, M., Balaprakash, P., Dorigo, M.: The ACO/F-RACE algorithm for combinatorial optimization under uncertainty. In Doerner, K. F. et. al, ed.: Metaheuristics - Progress in Complex Systems Optimization. Operations Research/Computer Science Interfaces Series, Springer Verlag, Berlin, Germany (2006) 189–203

[8] Balaprakash, P., Birattari, M., Stützle, T.: Improvement strategies for the F-Race algorithm: Sampling design and iterative refinement. In Bartz-Beielstein, T. et. al, ed.: Hybrid Metaheuristics. Volume 4771 of LNCS., Berlin, Germany, Springer-Verlag (2007) 108–122