

ACO/F-Race: Ant Colony Optimization and Racing Techniques for Combinatorial Optimization Under Uncertainty

Mauro Birattari*

Prasanna Balaprakash*

Marco Dorigo*

*IRIDIA, CP 194/6, Université Libre de Bruxelles
Avenue Franklin D. Roosevelt 50, 1050 Bruxelles, Belgium
{mbiro,pbalapra,mdorigo}@ulb.ac.be

Abstract

The paper introduces *ACO/F-Race*, an algorithm for tackling combinatorial optimization problems under uncertainty. The algorithm is based on *ant colony optimization* and on *F-Race*. The latter is a general method for the comparison of a number of candidates and for the selection of the best one according to a given criterion. Some experimental results on the PROBABILISTIC TRAVELING SALESMAN PROBLEM are presented.

1 Introduction

In a large number of real-world combinatorial optimization problems, the objective function is affected by uncertainty. In order to tackle these problems, it is customary to resort to a probabilistic model of the value of each feasible solution. In other words, a setting is considered in which the cost of each solution is a *random variable*, and the goal is to find the solution that minimizes some *statistics* of the latter. For a number of practical and theoretical reasons, it is customary to optimize with respect to the *expectation*. For a given probabilistic model, the expectation can always be computed but this typically involves particularly complex analytical manipulations and computationally expensive procedures. Two alternatives have been discussed in the literature: *analytical approximation* and *empirical estimation*. While the former explicitly relies on the underlying probabilistic model for approximating the expectation, the latter estimates the expectation through *sampling* or *simulation*.

In this paper we introduce *ACO/F-Race*, an *ant colony optimization* algorithm [7] for tackling combinatorial optimization problems under uncertainty with the *empirical estimation* approach. *F-Race* [5, 4] is an algorithm for the comparison of a number of candidates and for the selection of the best one. *F-Race* has been specially developed for tuning metaheuristics.¹ In the present paper, *F-Race* is used in an original way as a component of an *ant colony optimization* algorithm. More precisely, it is adopted for selecting the *best-so-far* ant, that is, the ant that is appointed for updating the pheromone matrix.

¹A public domain implementation of *F-Race* for R is available for download [3]. R is a language and environment for statistical computing that is freely available under the GNU GPL license.

The main advantage of the *estimation* approach over the one based on *approximation* is generality: Indeed, a sample estimate of the expected cost of a given solution can be simply obtained by averaging a number of realizations of the cost itself. On the other hand, computing a profitable approximation is a problem-specific issue and requires a deep understanding of the underlying probabilistic model. Since *ACO/F-Race* is based on the *empirical estimation* approach, it is straightforward to apply it to a large class of combinatorial optimization problems under uncertainty. For definiteness, in this paper we consider an application of *ACO/F-Race* to the PROBABILISTIC TRAVELING SALESMAN PROBLEM, more precisely to its *homogeneous* variant [9]. An instance of the PROBABILISTIC TRAVELING SALESMAN PROBLEM (PTSP) is defined as an instance of the well known TRAVELING SALESMAN PROBLEM (TSP), with the difference that in PTSP each city has a given probability of requiring being visited. In the paper we consider the *homogeneous* variant, in which the probability that a city must be visited is the same for all cities. PTSP is here tackled in the *a priori* optimization sense [1]: The goal is to find an *a priori* tour visiting all the cities, which minimizes the expected length of the associated *a posteriori* tour. The *a priori* tour must be found prior to knowing which cities indeed require being visited. The associated *a posteriori* tour is computed *after* knowing which cities need being visited, and is obtained by visiting them in the order in which they appear in the *a priori* tour. The cities that do not require being visited are simply skipped. This problem was selected as the first problem for testing the *ACO/F-Race* algorithm for two main reasons: First, PTSP is particularly simple to describe and to handle. In particular, the *homogeneous* variant is rather convenient since a single parameter, that is, the probability that each city requires being visited, defines the “stochastic character” of an instance: When the probability is one, we fall into the deterministic case; as it decreases, the normalized standard deviation of the cost of a given solution increases steadily. We can informally conclude that an instance of the *homogeneous* PTSP becomes *more and more stochastic* as the probability that cities require being visited decreases. This feature is particularly convenient in the analysis and visualization of experimental results. Second, some variants of *ant colony optimization* have been already applied to PTSP: Bianchi et al. [2] proposed *pACS*, a variant of *ant colony system* in which an *approximation by defect* of the expected length of the *a posteriori* tour is optimized. Gutjahr [8] proposed *S-ACO*, in which an *estimation* of the expected length of the *a posteriori* tour is optimized.

The rest of the paper is organized as follows: Section 2 discusses the problem of estimating, on the basis of a sample, the cost of a solution in a combinatorial optimization problem under uncertainty. Section 3 introduces the *ACO/F-Race* algorithm. Section 4 proposes some results obtained by *ACO/F-Race* on PTSP. Section 5 concludes the paper and highlights future research directions.

2 The empirical estimation of stochastic costs

For a formal definition of the class of problems that can be tackled by *ACO/F-Race*, we follow Gutjahr [8]:

$$\text{Minimize } F(x) = E[f(x, \omega)], \quad \text{subject to } x \in S, \quad (1)$$

where x is a solution, S is the set of feasible solutions, the operator E denotes the mathematical expectation, and f is the cost function which depends on x and also on a random (possibly

multivariate) variable ω . The presence of the latter makes the cost $f(x, \omega)$ of a given solution x a random variable.

In the *empirical estimation* approach to stochastic combinatorial optimization, the expectation $F(x)$ of the cost $f(x, \omega)$ for a given solution x is estimated on the basis of a sample $f(x, \omega_1), f(x, \omega_2), \dots, f(x, \omega_M)$, obtained from M independently-extracted realizations of the random variable ω :

$$\hat{F}_M(x) = \frac{1}{M} \sum_{i=1}^M f(x, \omega_i). \quad (2)$$

Clearly, $\hat{F}_M(x)$ is an *unbiased* estimator of $F(x)$.

In the case of PTSP, the elements of the general definition of a stochastic combinatorial optimization problem given above take the following meaning: A feasible solution x is an *a priori* tour visiting once and only once all cities. If cities are numbered from 1 to N , x is a permutation of $1, 2, \dots, N$. The random variable ω is extracted from an N -variate Bernoulli distribution and prescribes which cities need being visited. In the *homogeneous* variant of PTSP, each element in ω is independently extracted from a same univariate Bernoulli distribution with probability p , where p is a parameter defining the instance. The cost $f(x, \omega)$ is the length of an *a posteriori* tour visiting the cities indicated in ω , in the order in which they appear in x .

3 The ACO/F-Race algorithm

It is straightforward to extend an *ant colony optimization* algorithm for the solution, in the *empirical approximation* sense, of a combinatorial optimization problem under uncertainty. Indeed, it is sufficient to consider one single realization of the random influence ω , say ω' , and optimize the function $\hat{F}_1(x) = f(x, \omega')$. Indeed, $\hat{F}_1(x)$ is an *unbiased* estimator of $F(x)$. The risk we run by following this strategy is that we might sample a particularly *atypical* ω' which provides a misleading estimation of $F(x)$. A safer choice consists in considering a different realization of ω at each iteration of the *ant colony optimization* algorithm. The rationale behind this choice is that unfortunate modifications to the pheromone matrix that can be caused by sampling an *atypical* value of ω at a given iteration, will not have a large impact on the overall result and will be corrected in following iterations. In this paper we call *ACO-1* an *ant colony optimization* algorithm for stochastic problems in which the objective function is estimated on the basis of *one single* realization of ω which is sampled anew at each iteration of the algorithm.

A more refined approach has been proposed by Gutjahr [8] and consists in using a larger sample for estimating the value of $F(x)$. In Gutjahr's *S-ACO*, the solutions produced at a given iteration are compared on the basis of a single realization. The *iteration-best* is then compared with the *best-so-far* solution on the basis of a larger sample whose size is determined dynamically on the basis of a parametric statistical test: Further realizations are considered till when either a maximum amount of computation is performed, or when the difference between the sample means for the two solutions being compared is larger than 3 times their estimated standard deviation. The selected solution is stored as the new *best-so-far* for future comparisons and is used for updating the pheromone matrix.

The *ACO/F-Race* algorithm we propose in this paper is inspired by *S-ACO* and similarly to the latter it considers, at each iteration, a number of realizations for comparing candidate solutions and for selecting the best one which is eventually used for updating the pheromone matrix. The significant difference lays in the algorithm used at each iteration for selecting the best candidate solution. *ACO/F-Race* adopts *F-Race*, an algorithm originally developed for tuning metaheuristics [5, 4]. *F-Race* is itself inspired by a class of *racine* algorithms proposed in the machine learning literature for tackling the model selection problem [11, 12]. In *F-Race*, as in the other *racine* algorithms, a set of given candidates are sequentially evaluated on a number of test cases. As soon as sufficient evidence is gathered that a candidate is worse than at least another one, the former is discarded from the race and is not further evaluated. The race terminates when either one single candidate remains, or when a maximum amount of computation time is reached. The peculiarity of *F-Race* compared to other *racine* algorithms is the adoption of the *Friedman two-way analysis of variance by ranks* [6], a nonparametric statistical test that appears particularly suitable in the context of racing algorithms.

At each iteration of *ACO/F-Race*, m ants, where m is a parameter, construct a solution as it is customary in *ant colony optimization*. In particular, we have adopted here the *random-proportional* rule [7]. The m solutions built by the ants, together with the *best-so-far* solution, are evaluated and compared via *F-Race*. The procedure consists in a series of steps at each of which a new realization of ω is considered and is used for evaluating the solutions that are still in the race. At each step, a Friedman test is performed and solutions that are statistically dominated by at least another one are discarded from the race. The solution that *wins* the race is used for updating the pheromone and is stored as the *best-so-far* to be used in the following iteration of the algorithm.

4 Experimental results

In our experimental analysis we compare *ACO/F-Race* with *ACO-1* and *S-ACO*. The three algorithms differentiate only for what concerns the procedure used for selecting, at the end of each iteration, the *best-so-far* solution that gets reinforced. Apart for this element, the algorithms are identical. The implementation used in the experiments is based on Stützle [13]. The problems considered are *homogeneous* PTSP instances obtained from TSP instances generated by the DIMACS generator [10]. We present the results of two experiments. In the first, cities are *uniformly distributed*, in the second they are *clustered*. For each of the two experiments, we consider 100 TSP instances of 300 cities. Out of each TSP instance we obtain 21 PTSP instances by letting the probability range in $[0, 1]$ with a step of 0.05. The run time is of 60 seconds on an AMD Opteron™ 244. The three algorithms were not fine-tuned, and the parameters adopted are those suggested in Gutjahr [8] for *S-ACO*. This might possibly introduce a bias in favor of *S-ACO*. The solutions selected by each algorithm on each instance were then evaluated on 300 freshly-selected realizations. The results are reproduced in Figure 1.

Although trivial, the evaluation method adopted by *ACO-1* gives the best results when the variance of the length of the *a posteriori* tour is small (or null), that is, when the probability is close to 1. This is easily explained: using large samples for estimating $F(x)$ is simply a waste of time when the variance of $f(x, \omega)$ is low. On the other hand, when the probability




Figure 1: Experimental results on *uniformly distributed* and *clustered* instances. The plots represent the expected length of the *a posteriori* tour obtained by *ACO/F-Race* and *S-ACO*, normalized by the one obtained by *ACO-1*. The results proposed are obtained by running each algorithm on each instance for 60 seconds on an AMD Opteron™ 244.

decreases and the problem indeed becomes stochastic, considering a larger sample does repay, and *ACO/F-Race* and *S-ACO* obtain better results than *ACO-1*. Further, Figure 1 shows that *F-Race* improves over the parametric procedure adopted in *S-ACO*: Throughout the whole range of probabilities, *ACO/F-Race* is *significantly* better than *S-ACO* according to a paired Wilcoxon test ($\alpha = .01$).

5 Conclusions and future work

The preliminary experimental results proposed in Section 4 confirm the generality of the approach proposed in Gutjahr [8], and show that the *F-Race* algorithm can be profitably adopted for comparing solutions in the framework of applications of *ant colony optimization* to combinatorial optimization problems under uncertainty.

Further research is needed for properly assessing the quality of *ACO/F-Race*. We are currently developing an *estimation*-based local search for PTSP. We plan to study the behavior of *ACO/F-Race* enriched by this local search on *homogeneous* and *non-homogeneous* problems.

In the experimental analysis proposed in Section 4 the goal was to compare the evaluation procedure based on *F-Race* with the one proposed in Gutjahr [8] and with the trivial one based on a single sample. For this reason, solution construction and pheromone update were implemented as described in Gutjahr [8]. We plan to explore other possibilities, such as construction and update as defined in *MAX-MIN ant system* [14]. Applications to other problems, in particular of the VEHICLE ROUTING class, will be considered too.

Acknowledgments. This research has been supported by COMP²SYS, a Marie Curie Early Stage Research Training Site funded by the European Community's Sixth Framework Programme under contract number MEST-CT-2004-505079, and by the ANTS project, an *Action de Recherche Concertée* funded by the Scientific Research Directorate of the French Community of Belgium. Marco Dorigo acknowledges support from the Belgian FNRS, of which he is a Research Director.

Vienna, Austria, August 22–26, 2005

References

- [1] D. J. Bertsimas, P. Jaillet, and A. Odoni. A priori optimization. *Operations Research*, 38:1019–1033, 1990.
- [2] L. Bianchi, L. M. Gambardella, and M. Dorigo. Solving the homogeneous probabilistic travelling salesman problem by the ACO metaheuristic. In M. Dorigo, G. Di Caro, and M. Sampels, editors, *Ant Algorithms, 3rd International Workshop, ANTS 2002*, volume 2463 of *LNCS*, pages 176–187, Berlin, Germany, 2002. Springer-Verlag.
- [3] M. Birattari. Race. R package. <http://cran.r-project.org/>, 2003.
- [4] M. Birattari. *The Problem of Tuning Metaheuristics as Seen from a Machine Learning Perspective*. PhD thesis, Université Libre de Bruxelles, Brussels, Belgium, 2004.
- [5] M. Birattari, T. Stützle, L. Paquete, and K. Varrentapp. A racing algorithm for configuring metaheuristics. In W. B. Langdon, E. Cantú-Paz, K. Mathias, R. Roy, D. Davis, R. Poli, K. Balakrishnan, V. Honavar, G. Rudolph, J. Wegener, L. Bull, M. A. Potter, A. C. Schultz, J. F. Miller, E. Burke, and N. Jonoska, editors, *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 11–18, San Francisco, CA, USA, 2002. Morgan Kaufmann.
- [6] W. J. Conover. *Practical Nonparametric Statistics*. John Wiley & Sons, New York, NY, USA, third edition, 1999.
- [7] M. Dorigo and T. Stützle. *Ant Colony Optimization*. MIT Press, Cambridge, MA, USA, 2004.
- [8] W. J. Gutjahr. S-ACO: An ant based approach to combinatorial optimization under uncertainty. In M. Dorigo, M. Birattari, C. Blum, L. M. Gambardella, F. Mondada, and T. Stützle, editors, *Ant Colony Optimization and Swarm Intelligence, 4th International Workshop, ANTS 2004*, volume 3172 of *LNCS*, pages 238–249, Berlin, Germany, 2004. Springer-Verlag.
- [9] P. Jaillet. *Probabilistic Travelling Salesman Problems*. PhD thesis, The Massachusetts Institute of Technology, Cambridge, MA, USA, 1995.
- [10] D. S. Johnson, L. A. McGeoch, C. Rego, and F. Glover. 8th DIMACS implementation challenge. <http://www.research.att.com/~dsj/chtsp/>, 2001.
- [11] O. Maron and A. W. Moore. Hoeffding races: Accelerating model selection search for classification and function approximation. In J. D. Cowan, G. Tesauro, and J. Alspector, editors, *Advances in Neural Information Processing Systems*, volume 6, pages 59–66, San Francisco, CA, USA, 1994. Morgan Kaufmann.
- [12] A. W. Moore and M. S. Lee. Efficient algorithms for minimizing cross validation error. In *Proceedings of the Eleventh International Conference on Machine Learning*, pages 190–198, San Francisco, CA, USA, 1994. Morgan Kaufmann.
- [13] T. Stützle. ACOTSP, version 1.0. <http://aco-metaheuristic.org/>, 2002.
- [14] T. Stützle and H. H. Hoos. *MAX-MIN* ant system. *Future Generation Computer Systems*, 16(8):889–914, 2000.

Vienna, Austria, August 22–26, 2005