Theoretical Computer Science 540–541 (2014) 156–168

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

On the computational complexity of the Probabilistic Traveling Salesman Problem with Deadlines

D. Weyland

IDSIA – Dalle Molle Institute for Artificial Intelligence, USI/SUPSI, Galleria 2, 6928 Manno-Lugano, Switzerland

ARTICLE INFO

Keywords:
Stochastic combinatorial optimization
Stochastic vehicle routing
Computational complexity
Inapproximability

ABSTRACT

The Probabilistic Traveling Salesman Problem with Deadlines (PTSPD) is a stochastic vehicle routing problem considering time dependencies in terms of deadlines. The evaluation of the PTSPD objective function is believed to be a computationally demanding task and all efforts to find a polynomial time approach have failed so far. On the other hand, no hardness results regarding the evaluation of the PTSPD objective function are known. We fill this gap and show that the evaluation of the PTSPD objective function is indeed a computationally demanding task: even for Euclidean instances this task is #P-hard. We then derive additional hardness results for different computational tasks related to the PTSPD. Among other results, we show that the decision variant and the optimization variant of the PTSPD are both #P-hard. Following this, we focus on polynomial time approximations of the PTSPD objective function. We prove the existence of an FPRAS under certain conditions and examine the approximation guarantees obtained by the existing approaches. Finally, we give a strong inapproximability result regarding the objective function of a slightly more general problem, the Dependent PTSPD.

1. Introduction

Stochastic combinatorial optimization problems have received increasing attention in recent years. By using stochastic information as input data, more realistic models of real world problems can be obtained. As a trade-off, such problems are usually much harder to solve than their non-stochastic counterparts and it is a great challenge to develop efficient solution approaches for these problems.

An important group of stochastic combinatorial optimization problems are the so called stochastic vehicle routing problems. These are problems arising in the field of transportation and logistics for which the input is (partially) modeled in a stochastic way. Many different stochastic vehicle routing problems have been investigated. Among the most typical stochastic vehicle routing problems are the Probabilistic Traveling Salesman Problem (PTSP, [15,5,2]), the Vehicle Routing Problem with Stochastic Demands [19,8] and the Vehicle Routing Problem with Stochastic Demands and Customers [14,13,12]. Although time dependencies are widely used in models for non-stochastic vehicle routing problems [18,9,1,10,22,23], only a few stochastic vehicle routing problems consider such time dependencies. One of them is the recently introduced Probabilistic Traveling Salesman Problem with Deadlines (PTSPD, [6]). This problem is a generalization of the PTSP, where in addition time dependencies are modeled by means of deadlines. A detailed introduction, including a thorough motivation for this problem, is given in [6,7].

While the objective function for the PTSP can be evaluated in polynomial time, it is believed that this is a computationally demanding task for the PTSPD. Until now no hardness results were known regarding the objective function of the

E-mail address: dennis@idsia.ch.

0304-3975/$ – see front matter
http://dx.doi.org/10.1016/j.tcs.2013.11.028

PTSPD. In this work we show that computing the probabilities with which deadlines are violated is already #P-hard for Euclidean instances of the PTSPD. Starting from this result we derive hardness results for various computational tasks related to the PTSPD. Among other results, we show that the evaluation of the PTSPD objective function, the optimization variant of the PTSPD and the decision variant of the PTSPD are #P-hard, even for Euclidean instances. The former results are the first of this kind and the latter results strengthen the trivial hardness results that the PTSPD inherits as a generalization of the well-known Traveling Salesman Problem.

An important open problem in this context is the approximability of the PTSPD objective function. In our work we derived several results that could help solving this issue. We first show that there exists a fully polynomial randomized approximation scheme (FPRAS) for an important restricted case of the PTSPD. We then continue with an investigation of the approximation guarantees of the known approximation approaches for the PTSPD objective function. Here we show that none of these approaches can guarantee a satisfactory worst-case approximation ratio in the general case and most of them also fail in the more restricted case mentioned before. After that we focus on a generalization of the PTSPD, in which certain dependencies among the stochastic input are allowed. We are able to give a strong inapproximability results for the objective function of this slightly more general problem, the so-called Dependent PTSPD.

The remaining part of this article is organized in the following way. In Section 2 we introduce the PTSPD. We give a formal description of the problem and its objective function together with an overview about the existing literature related to this problem. Section 3 contains the results about the computational complexity of several computational tasks related to the PTSPD. In Section 4 we then investigate the approximability of the PTSPD objective function. Section 5 deals with the inapproximability results for the objective function of the Dependent PTSPD. Finally, we finish the paper with a discussion of the results and with an outlook for possible further research in Section 6.

Please note that this article is an extension of the conference paper *Hardness Results for the Probabilistic Traveling Salesman Problem* [26]. While the results in Section 3 are based on the conference paper, the results of Sections 4 and 5 are completely new. The new results are of great importance for the Probabilistic Traveling Salesman Problem with Deadlines and can be generalized to other stochastic vehicle routing problems and other stochastic combinatorial optimization problems. Additionally, we could solve some of the open problems raised in [26].

2. The Probabilistic Traveling Salesman Problem with Deadlines

The Probabilistic Traveling Salesman Problem with Deadlines (PTSPD) has been introduced in 2008 [6]. Four different models are proposed: three recourse models [4] and one chance constrained model [4]. A representation of the chance constrained model as a linear program and computational approaches for the evaluation of the objective function for the recourse models are introduced. The computational time for the evaluation of solutions under the recourse models is not bounded by a polynomial in general and it seems that the evaluation of the objective function is a computationally demanding task for all the models. So far there is no polynomial time approach known. Additionally, special cases in which the problem can be solved efficiently, or in which at least the objective function can be evaluated efficiently, are discussed thoroughly. Due to the computationally demanding objective function three different approximations for this objective function have been introduced in [7]. Two of those approximations achieve a polynomial runtime (in case appropriate parameters are used). Moreover, a series of computational experiments using one of the recourse models were performed. The results show that a huge amount of the computational time for certain heuristics can be saved if the approximations of the objective function are used instead of the exact objective function.

In our work we focus on one of the recourse models from [6] (Recourse I, fixed penalties). Note that our results also hold for the other variants. The proofs for the other variants are based on the same ideas and require only slight modifications. In fact, the proofs for the chance constrained model are much simpler. The PTSPD (Recourse I, fixed penalties) can be defined as follows. We have given the location of a depot, the location of different customers and travel times between all of them. Note that in the context of this problem the terms travel times and distances are used interchangeably, assuming unit velocity. For each customer we have given a probability, with which this customer requires to be visited. Additionally, for each customer we have given a deadline, until which this customer should have been visited, as well as a penalty value, which represents the fixed costs for a missed deadline. A solution for this problem is represented by a tour starting at the depot, visiting all of the customers exactly once and finishing at the depot. In this context such a solution is called an a-priori solution. For given realizations of the random events, an a-posteriori solution is derived from this a-priori solution. Here customers that do not require to be visited are just skipped, while the other customers are still visited in the order defined by the a-priori solution. The costs for such an a-posteriori solution are the total travel times plus the penalties for missed deadlines. The optimization goal is to find an a-priori solution minimizing the expected costs over the corresponding a-posteriori solutions under the given probability distributions.

More formal, the problem can be defined in the following way. We have given a set V with a special element $v_0 \in V$. We call v_0 the depot and $V \setminus \{v_0\}$ the set of customers. Additionally, we have given a function $d : V \times V \to \mathbb{N}$ representing the travel times, a function $p : V \setminus \{v_0\} \to [0,1]$ representing the probabilities with which customers require a visit, a function $t : V \setminus \{v_0\} \to \mathbb{N}$ representing the deadlines of the customers and a function $h : V \setminus \{v_0\} \to \mathbb{N}$ representing the fixed penalties for each customer. An a-priori solution is then represented by a permutation $\tau : \{1, 2, \ldots, |V|\} \to V$ with $\tau(1) = v_0$. Now the optimization goal is to find an a-priori solution such that the expected costs over the corresponding a-posteriori solutions