Computers & Operations Research 40 (2013) 1716–1728

Contents lists available at SciVerse ScienceDirect

Computers & Operations Research

journal homepage: www.elsevier.com/locate/caor

A memetic algorithm for the travelling salesperson problem with hotel selection

Marco Castro [a,*], Kenneth Sörensen [a], Pieter Vansteenwegen [b,c], Peter Goos [a,d]

[a] University of Antwerp, Belgium
[b] Ghent University, Belgium
[c] KU Leuven, Belgium
[d] Erasmus University Rotterdam, The Netherlands

ARTICLE INFO

Available online 16 January 2013

Keywords:
Memetic algorithm
TSP
Hotel selection

ABSTRACT

In this paper, a metaheuristic solution procedure for the travelling salesperson problem with hotel selection (TSPHS) is presented. The metaheuristic consists of a memetic algorithm with an embedded tabu search, using a combination of well-known and problem-specific neighbourhoods. This solution procedure clearly outperforms the only other existing metaheuristic in the literature. For smaller instances, whose optimal solution is known, it is able to consistently find the optimal solution. For the other instances, it obtains several new best known solutions.

1. Introduction

The *travelling salesperson problem with hotel selection* (TSPHS) was recently introduced by Vansteenwegen et al. [37]. The motivation for this problem is that a salesperson often cannot visit all customers in a single day, due to the fact that he/she can only work for a limited number of hours per day. This implies that the salesperson needs to select a hotel each night, on top of determining the optimal sequence in which to visit all customers. Every day should start and end in one of the available hotels and, if a given day ends in a certain hotel, the next day should start in the same hotel. The primary goal of this problem is to minimise the required number of days, while the secondary goal is to minimise the total travelled length.

Throughout the paper, the term "trip" is used to indicate a sequence of customers, starting and ending in a hotel, while the term "tour" is used for a complete sequence of connected trips that, together, visits all customers.

Although this problem appears to be very similar to the (regular) travelling salesperson problem (TSP), it is inherently more difficult due to the hotel selection requirement. The selected hotels determine to a large extent the length of the total tour.

A number of applications of the TSPHS are presented in Vansteenwegen et al. [37]: the travelling salesperson who needs several days to visit all customers; a multi-day trip for a truck driver in which every day trip should start and end on an appropriate parking space; a multi-day tourist visit to a certain region; mailmen who want to split their round into a number of connected sub-rounds in order to lighten their bag. Furthermore, the TSPHS may be used to route electric vehicles in which routes are split into trips whose maximum duration is constrained by the battery charge, and batteries can be swapped or recharged at intermediate points.

The rest of the paper is organised in the following way. In Section 2, a review of the relevant literature is presented. In Section 3, a description of the problem and a modified MIP formulation is introduced, while, in Section 4, a metaheuristic procedure for the TSPHS is outlined. In Section 5, the experiments as well as a parametric analysis are presented. Finally, Section 6 provides a conclusion and some suggestions for future research.

2. Literature review

Several problems related to the TSPHS can be found in the literature. In the multiple travelling salesperson problem (mTSP) [5], a number of salespeople, all starting and ending in the same depot, are available to visit all customers. In the vehicle routing problem (VRP) [36], the objective is to minimise the total distance travelled by a number of vehicles, each limited by a given capacity. Contrary to the TSPHS, both problems only consider one depot (or hotel). The constraint that all vehicles should start and end at a single depot is relaxed in the multi-depot vehicle routing problem (MDVRP) [9,30] where several depots are available, each with a fleet of vehicles. However, in the MDVRP, each

* Corresponding author. Tel.: +32 3 265 4061; fax: +32 3 265 4901.
E-mail addresses: marco.castro@ua.ac.be (M. Castro), kenneth.sorensen@ua.ac.be (K. Sörensen), pieter.vansteenwegen@cib.kuleuven.be (P. Vansteenwegen), peter.goos@ua.ac.be (P. Goos).

0305-0548/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.cor.2013.01.006

vehicle must start and end at the same depot, while, in the TSPHS, a trip may start at one hotel and end at another hotel. In location-routing problems (LRP) [28], the depots are not fixed in advance. In the basic LRP, a number of depots has to be selected from a given set, in order to minimise the total cost of using the selected depots and routing the vehicles starting from these depots. Each vehicle must return to the depot it started from. The mTSP, VRP, MDVRP and the LRP thus each have some features in common with the TSPHS, but the most important differences are that, in the TSPHS, only one vehicle or salesperson is available and all trips need to be connected.

In the context of problems with intermediate facilities (IFs), several problems related to hotel selection arise in the literature. In the periodic vehicle routing problem with intermediate facilities (PVRP-IF) [1], a depot is fixed in advance; customers are served during a work shift whose maximum duration cannot exceed an established time limit; and the vehicle may be replenished at one of the available intermediate facilities. In the waste collection vehicle routing problem with time windows (WCVRPTW) [6,22] a depot, a set of customers and a set of waste disposal facilities are available. The empty vehicles depart from the depot, collect the waste from a set of customers and are emptied at the disposal facilities. The vehicles may visit the disposal facilities as many times as needed and, at the end of the work shift, the vehicles return to the depot emptily. In the multi-depot vehicle routing problem with inter-depot routes (MDVRPI) [11], a number of depots is available and routes are designed to serve the customers. The routes may start and end at the same depot or at different depots, while the time needed by each vehicle to traverse the set of routes assigned to it stays within a certain time limit. A variant of this problem is the vehicle routing problem with intermediate replenishment facilities (VRPIRF) [35] in which a fleet of vehicles is based at a single central depot.

Several arc routing problems involving intermediate facilities also exhibit similarities with the TSPHS. In the capacitated arc routing problem with intermediate facilities (CARPIF) [17,29], an extension of the capacitated arc routing problem (CARP) [20], a set of edges in a graph represent a road network. A demand and a travel time are associated with each edge. A subset of nodes in the graph, referred to as the intermediate facilities (IFs), represents available replenishment facilities. A fleet of vehicles with homogeneous capacity is available at a central depot. Loaded vehicles depart from the depot, traverse the edges on the graph servicing demands and may be replenished at one of the available IFs. The objective is to determine the set of vehicle routes that minimises the total travelled time. Several real-world applications of this problem are mentioned in Polacek et al. [29]. The arc routing problem with intermediate facilities under capacity and length restrictions (CLARPIF) is presented as a variant of the CARPIF in Ghiani et al. [16]. In this problem, an upper bound is imposed on the length of each route.

Just like the TSPHS, which allows each hotel to be visited on multiple occasions, the PVRP-IF, WCVRPTW, MDVRPI, VRPIRF, CARPIF and CLARPIF allow multiple visits to intermediate facilities. One difference between these problems and the TSPHS is that there is no explicit upper bound on the trip length. Instead, the trip lengths are indirectly bounded by the capacity of the vehicle(s). It is true that some of the VRPs and ARPs involve an upper bound on the travel time, but this upper bound only applies to the total travel time. This is unlike the TSPHS, where the upper bound on the travel time applies to each day trip. Another important difference between the VRPs and ARPs, on the one hand, and the TSPHS, on the other hand, is that the primary objective of the TSPHS is the minimisation of the number of trips (equivalent to the number of visits to intermediate facilities), rather than minimising the duration of the entire tour.

3. Problem description

Given a non-empty set H of hotels, and a set C of customers, the TSPHS is defined on a complete graph $G=(V,A)$ where $V=H\cup C$ and $A=\{(i,j)|i,j\in V, i\neq j\}$. Each customer $i\in C$ requires a service or visiting time τ_i (with $\tau_i=0, \forall i\in H$). The time c_{ij} needed to travel from location i to j is known for all pairs of locations. The goal is to first minimise the number of connected trips required to visit all customers, and then to minimise the total travel time of the tour. The initial and final hotel of the tour (i.e., the starting point of the first trip and the end point of the final trip) are assumed to be identical and given ($i=0$). This hotel can also be used as an intermediate hotel during the tour. Furthermore, (a) each trip must start and end in one of the available hotels, (b) the travel time of each trip must not exceed a given time budget L, and (c) a trip should start in the hotel where the previous trip ended. Since there is no limit on the number of visits to a hotel, a solution to the TSPHS is not necessarily a single cycle [37].

In this paper, an IP formulation is presented that modifies the formulation of Vansteenwegen et al. [37] in two ways: (1) a weighted objective function is used to circumvent the non-linearity problem that results from the (lexicographical) ordering of the two objectives and (2) the Miller–Tucker–Zemlin sub-tour elimination constraints were replaced by the much more efficient Dantzig–Fulkerson–Johnson constraints, allowing the solver to find more optimal solutions in a smaller computing time. In order to prioritise the minimisation of the number of trips, the number of trips is multiplied by a sufficiently large number M in the objective function so that a solution involving a smaller number of trips always has a better objective function value than any other solution involving a larger number of trips.

Let x_{ij}^d be a binary variable that takes the value 1 if, in trip d, a visit to a customer or hotel i is followed by a visit to customer or hotel j, and the value 0 otherwise. Also, let the binary variable y^d take the value 1 if, in trip d, at least one customer or hotel is visited, and the value 0 otherwise. Thus, y^d will be zero if no trip is required on day d in order to visit all customers. In the IP formulation's objective function, the variables y^d and x_{ij}^d are used for the mathematical expression of the primary and secondary objective, respectively. Finally, let D be the maximum number of trips contained in the solution.

$$\min \quad M\sum_{d=1}^{D} y^d + \sum_{d=1}^{D}\left(\sum_{(i,j)\in A} c_{ij}x_{ij}^d\right) \tag{1}$$

$$\text{s. t.} \quad \sum_{d=1}^{D}\sum_{i\in V} x_{ij}^d = 1, \quad j\in C \tag{2}$$

$$\sum_{i\in V} x_{ij}^d = \sum_{i\in V} x_{ji}^d, \quad j\in C,\ d=1,\ldots,D \tag{3}$$

$$\sum_{h\in H}\sum_{j\in V\setminus\{h\}} x_{hj}^d = y^d, \quad d=1,\ldots,D \tag{4}$$

$$\sum_{h\in H}\sum_{i\in V\setminus\{h\}} x_{ih}^d = y^d, \quad d=1,\ldots,D \tag{5}$$

$$\sum_{(i,j)\in A} (c_{ij}+\tau_j)x_{ij}^d \le L, \quad d=1,\ldots,D \tag{6}$$

$$\sum_{j\in V\setminus\{0\}} x_{0j}^1 = 1 \tag{7}$$

$$\sum_{i\in V\setminus\{0\}} x_{i0}^d \ge y^d - y^{d+1}, \quad d=1,\ldots,D-1 \tag{8}$$