European Journal of Operational Research 255 (2016) 422–441

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Production, Manufacturing and Logistics

Large neighborhood search for multi-trip vehicle routing

Véronique François[a], Yasemin Arda[a,*], Yves Crama[a], Gilbert Laporte[b]

[a] QuantOM, HEC Management School, University of Liège, Belgium
[b] Canada Research Chair in Distribution Management, HEC Montréal, Canada

ARTICLE INFO

Article history:
Received 27 November 2015
Accepted 30 April 2016
Available online 10 May 2016

Keywords:
Transportation
Vehicle routing
Multi-trip
Large neighborhood search
Automatic algorithm configuration

ABSTRACT

We consider the multi-trip vehicle routing problem, in which each vehicle can perform several routes during the same working shift to serve a set of customers. The problem arises when customers are close to each other or when their demands are large. A common approach consists of solving this problem by combining vehicle routing heuristics with bin packing routines in order to assign routes to vehicles. We compare this approach with a heuristic that makes use of specific operators designed to tackle the routing and the assignment aspects of the problem simultaneously. Two large neighborhood search heuristics are proposed to perform the comparison. We provide insights into the configuration of the proposed algorithms by analyzing the behavior of several of their components. In particular, we question the impact of the roulette wheel mechanism. We also observe that guiding the search with an objective function designed for the multi-trip case is crucial even when exploring the solution space of the vehicle routing problem. We provide several best known solutions for benchmark instances.

1. Introduction

The multi-trip vehicle routing problem (MTVRP) is a variant of the vehicle routing problem (VRP) in which each vehicle can perform one or more trips, also called *routes*, during the same planning period. The set of routes performed by a given vehicle constitutes a *tour* whose total duration cannot exceed a given time limit. The MTVRP arises when the demands of the customers are large compared to vehicle capacities or when the distances are relatively short. In food distribution systems, drivers sometimes perform two, three or even more delivery routes during the same working day.

It is common to solve the MTVRP by combining VRP and bin packing (BP) algorithms (Fleischmann, 1990; Olivera & Viera, 2007; Petch & Salhi, 2004; Salhi & Petch, 2007; Taillard, Laporte, & Gendreau, 1996). Vehicle routes are first obtained by applying VRP algorithms. These routes are then assigned to a fleet with a limited number of vehicles, generally by applying BP techniques: each route is viewed as an item whose size corresponds to its duration and each vehicle as a bin of capacity equal to the maximum allowed tour duration.

The main contribution of this work is to propose local search operators specifically designed for multi-trip variants of the VRP. They stem from classical VRP operators and take into consideration the routing and route assignment aspects of multi-trip problems. Our aim is to compare the performance of algorithms that incorporate these multi-trip operators with those that treat the routing and the packing subproblems separately. Two adaptive large neighborhood search (ALNS) algorithms (Pisinger & Ropke, 2007; Ropke & Pisinger, 2006a,b) are described: the ALNS with multi-trip operators (ALNSM) and the ALNS combined with BP (ALNSP). Both algorithms are tested on the benchmark instances of Taillard et al. (1996). The ALNSM yields very good results but is outperformed by the ALNSP which produces 10 new best known MTVRP solutions.

As a second contribution, we analyze the behavior of various ALNS components and we describe the interactions between some of these. Every implementation option that was considered during the design phase of the algorithms is given, not only the most efficient one, thus providing insights into the global behavior of the proposed ALNS metaheuristics. We use the `irace` package (López-Ibáñez, Dubois-Lacoste, Stützle, & Birattari, 2011), an automatic configuration tool, not only as a fine-tuning engine, but also as a means to gain meaningful algorithmic insights.

Section 3 describes the MTVRP along with some notations. Section 4 presents the specific multi-trip operators. Section 5 details the ALNSM and ALNSP implementations. The results are then presented in Section 6, along with further experiments about algorithm components in Section 6.4. Section 7 presents the conclusions.

* Corresponding author.
E-mail addresses: veronique.francois@ulg.ac.be (V. François), yasemin.arda@ulg.ac.be (Y. Arda), yves.crama@ulg.ac.be (Y. Crama), gilbert.laporte@cirrelt.ca (G. Laporte).

http://dx.doi.org/10.1016/j.ejor.2016.04.065
0377-2217/

2. Literature review

Fleischmann (1990) was the first to address the MTVRP. He combined a modified savings heuristic with a BP heuristic. Taillard et al. (1996) later proposed a three-phase method. First, a large number of VRP routes are generated by constructing and modifying VRP solutions using a tabu search (TS) algorithm. Several VRP solutions are then built based on the routes constructed in this first step. Finally, a BP procedure is applied to each VRP solution in an attempt to generate feasible MTVRP solutions. In Brandão and Mercer (1998), a constructive procedure sequentially builds an MTVRP solution that may involve overtime. A TS procedure then refines this solution by applying classical VRP operators. Petch and Salhi (2004) proposed a multi-phase construction algorithm in which VRP solutions are iteratively created and refined. Each is tentatively transformed through a BP procedure until a feasible MTVRP solution is found or a stopping criterion is satisfied. Salhi and Petch (2007) applied a genetic algorithm in which chromosomes represent ordered circular sectors. A savings heuristic is applied to solve VRPs within the circular sectors, and a BP heuristic then creates MTVRP solutions out of the routes produced by the savings procedure. Olivera and Viera (2007) proposed an adaptive memory programing (AMP) procedure combining TS and a BP heuristic. A memory of routes is initialized by constructing VRP solutions by means of the sweep algorithm. The algorithm then iteratively creates and improves new VRP solutions from routes that are randomly chosen in the memory. These VRP solutions are improved by applying a TS procedure. A BP heuristic is used at each iteration of the TS heuristic in order to tentatively produce a feasible MTVRP solution. The new routes are added to the adaptive memory and sorted according to the quality of the MTVRP solution to which they belong. Cattaruzza, Absi, Feillet, and Vidal (2014) provided state-of-the-art results for the MTVRP. These authors proposed a memetic algorithm in which each chromosome defines a customer sequence. A modified split procedure (Prins, 2004) partitions the customer sequences to tentatively obtain an MTVRP solution. The authors described a second version of their algorithm in which a local search procedure reassigns routes to vehicles while at the same time performing a VRP move. Recently, Mingozzi, Roberti, and Toth (2013) developed an exact algorithm for the MTVRP which yielded an optimal solution for 42 out of the 106 benchmark instances of Taillard et al. (1996). Their model combines a partitioning formulation with valid inequalities. For an extensive review of MTVRP variants, the interested reader is referred to Cattaruzza, Absi, and Feillet (2016).

3. Problem description and notations

Let $G(V, E)$ be a complete undirected graph where $V = \{0, \ldots, n\}$ is the set of nodes and $E = \{(i, j) | i, j \in V, i < j\}$ is the set of edges. Each node $i = 1, \ldots, n$ represents a customer, while node $i = 0$ represents the depot. With each customer $i = 1, \ldots, n$ is associated a demand d_i that must be satisfied by exactly one delivery (i.e., split deliveries are not allowed). A fleet of m identical vehicles is based at the depot. The travel time on edge $(i, j) \in E$ is t_{ij}. Each vehicle $k = 1, \ldots, m$ has a limited capacity Q and a maximum allowed working duration T_{max}, and must perform a tour $\mathcal{T}_k$ made up of a set of routes starting and ending at the depot. The total demand of the customers served by any route of $\mathcal{T}_k$ must not exceed Q, and the time needed to perform $\mathcal{T}_k$ must not exceed T_{max}. The objective is to determine a set of tours minimizing the total travel time while satisfying the constraints.

The metaheuristics that we have developed work on a relaxed version of the MTVRP called the R-MTVRP, in which the tour duration constraints are not considered. In the following, we denote MTVRP solutions by $\mathcal{S}$, and R-MTVRP solutions by $\hat{\mathcal{S}}$. If at least one vehicle of an R-MTVRP solution $\hat{\mathcal{S}}$ travels for a duration that exceeds T_{max}, then $\hat{\mathcal{S}}$ contains some overtime. The overtime of vehicle k, denoted by O_k, is defined as $O_k = \max\{0, D_k - T_{max}\}$, where D_k is the total duration of tour $\mathcal{T}_k$. The total overtime of $\hat{\mathcal{S}}$ is defined as $O_{\hat{\mathcal{S}}} = \sum_{k=1,\ldots,m} O_k$. The objective function of the R-MTVRP includes the total travel time, as well as the penalized overtime. The associated penalization factor will be introduced later on. R-MTVRP solutions that do not contain overtime are also feasible MTVRP solutions.

Solutions to the well-known capacitated vehicle routing problem (CVRP) are used to explore the solution space during the execution of the ALNSP heuristic. The CVRP is defined on the same graph $G(V, E)$ as the MTVRP. The fleet size is unlimited, while each vehicle having a capacity Q can perform only one route. The objective function to be minimized is the total travel time. No cost is incurred for using the vehicles, i.e., the number of routes has no impact on the objective function. Let $\mathcal{X}$ be the set of routes that constitute a feasible CVRP solution. In order to transform $\mathcal{X}$ into an R-MTVRP solution, each route of $\mathcal{X}$ has to be assigned to one of the m available vehicles. All assignments of the routes in $\mathcal{X}$ that satisfy the MTVRP tour duration constraints yield MTVRP solutions that are equivalent in terms of their objective function value since they have the same duration and contain no overtime. Finding such a feasible assignment is equivalent to solving a BP problem where each route is an item with a size corresponding to its duration, and each vehicle is a bin of capacity equal to the maximum allowed tour duration. Since the CVRP objective function is simply the total travel time, if $\mathcal{X}$ is an optimal CVRP solution, then each feasible assignment of its routes provides an optimal MTVRP solution.

Table 1 presents the recurring notations of this paper.

4. Specific local search operators

The purpose of the proposed multi-trip operators is to manage the routing and the assignment aspects of the problem simultaneously, instead of treating them independently. In MTVRP or R-MTVRP solutions, any reordering of the routes of a given tour $\mathcal{T}_k$ leaves its total duration and overtime unchanged. However, the operators presented below are not designed to treat routes as separate entities. Instead, they treat any tour $\mathcal{T}_k$ as a giant tour made up of the routes of vehicle k, i.e., an ordered sequence of routes. More precisely, the representation of a tour starts with an origin depot, ends with a destination depot, and contains the customer sequence of each of its routes. Each of these customer sequences is separated by a depot that we call an “internal depot”. In the following, routes in $\mathcal{T}_k$ are said to be consecutive if their respective customer sequences are only separated by one internal depot in the giant tour representation of $\mathcal{T}_k$.

4.1. Removal and insertion operators

The specific removal and insertion operators presented below adapt to the MTVRP context destroy-and-repair heuristics (similar to those used by Ropke and Pisinger (2006a) and Pisinger and Ropke (2007)).

4.1.1. Inserting a customer in a multi-trip context

Let $\hat{\mathcal{S}}$ be a partial solution of a given R-MTVRP. In $\hat{\mathcal{S}}$, a sequence of nodes (depots and customers) forming a tour $\mathcal{T}_k$ is known for each available vehicle k, but some of the customers are not assigned to any vehicle. Some of the vehicles may also be empty (the representation of the associated tour contains only the origin and the destination depots). The solution $\hat{\mathcal{S}}$ needs to be repaired by inserting in it the unrouted customers. Let v_i and v_{i+1} be two consecutive nodes of a tour $\mathcal{T}_k = (\ldots, v_i, v_{i+1}, \ldots)$ in $\hat{\mathcal{S}}$. We consider