Computers & Operations Research 58 (2015) 87–99

Contents lists available at ScienceDirect

Computers & Operations Research

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

Hybrid metaheuristics for the Clustered Vehicle Routing Problem

Thibaut Vidal [a], Maria Battarra [b,*], Anand Subramanian [c], Güneş Erdoğan [d]

[a] Massachusetts Institute of Technology, Laboratory for Information and Decision Systems, 77 Massachusetts Avenue, Cambridge, MA 02139, USA
[b] University of Southampton, Mathematical Sciences, Southampton SO17 1BJ, UK
[c] Universidade Federal da Paraíba, Departamento de Engenharia de Produção, Centro de Tecnologia, Campus I - Bloco G, Cidade Universitária, João Pessoa-PB, 58051-970, Brazil
[d] University of Bath, School of Management, Bath, BA2 7AY, UK

ARTICLE INFO

Available online 13 November 2014

Keywords:
Clustered Vehicle Routing
Iterated Local Search
Hybrid Genetic Algorithm
Large neighborhoods
Shortest path

ABSTRACT

The Clustered Vehicle Routing Problem (CluVRP) is a variant of the Capacitated Vehicle Routing Problem in which customers are grouped into clusters. Each cluster has to be visited once, and a vehicle entering a cluster cannot leave it until all customers have been visited. This paper presents two alternative hybrid metaheuristic algorithms for the CluVRP. The first algorithm is based on an Iterated Local Search algorithm, in which only feasible solutions are explored and problem-specific local search moves are utilized. The second algorithm is a hybrid genetic search, for which the shortest Hamiltonian path between each pair of vertices within each cluster should be precomputed. Using this information, a sequence of clusters can be used as a solution representation and large neighborhoods can be efficiently explored, by means of bi-directional dynamic programming, sequence concatenation, and appropriate data structures. Extensive computational experiments are performed on benchmark instances from the literature, as well as new large scale instances. Recommendations on the choice of algorithm are provided, based on average cluster size.

1. Introduction

This paper addresses the *Clustered Vehicle Routing Problem* (CluVRP), which has been introduced by Sevaux and Sörensen [17]. The CluVRP is defined over an undirected graph $\mathcal{G}=(\mathcal{V},\mathcal{E})$, where the vertex 0 is the *depot* and any other vertex $i \in \mathcal{V} \setminus \{0\}$ is a customer with demand $q_i > 0$. A fleet of m vehicles, each with capacity Q, is stationed at the depot. The set of customers is partitioned into N disjoint and nonempty subsets called *clusters*, such that $\mathcal{V} = V_1 \cup \cdots \cup V_N$. The customers in each cluster have to be visited consecutively, that is, the vehicle visiting a customer in the cluster cannot leave the cluster until all the other customers in the cluster have been visited. Each edge $(i,j) \in \mathcal{E}$ is associated with a travel cost c_{ij}, and the objective is to minimize the total travel cost. The CluVRP is a generalization of the Capacitated Vehicle Routing Problem (CVRP, cf. book of [24]), obtained when each cluster contains a single vertex, and of the Clustered Traveling Salesman Problem (CluTSP, [5]), obtained when $m=1$. The CVRP and the CluTSP are both $\mathcal{NP}$−Hard, and so is the CluVRP.

Sevaux and Sörensen [17] introduced the CluVRP in the context of a real-world application where containers are employed to carry goods. The customers expecting parcels in the same container form a cluster, because the courier has to deliver the content of a whole container before handling another container. Clusters also arise in applications involving passenger transportation, where passengers prefer to travel with friends or neighbors (as in the transportation of elderly to recreation centers). Gated communities (residential or industrial areas enclosed in walled enclaves for safety and protection reasons) provide another natural example of clusters. The customers within a gated community are likely to be visited by a single vehicle in a sequence, otherwise the vehicles have to spend additional time for the security controls at the gates.

Clusters can thus be imposed by the geography, the nature of the application, as well as by practitioners aiming to achieve *compact* and easy-to-implement routing solutions. Clustered routes allow drivers to be assigned to areas (i.e., certain streets

* Corresponding author. Tel.: +44 23 8059 3863.
E-mail addresses: vidal@mit.edu (T. Vidal), m.battarra@soton.ac.uk (M. Battarra), anand@ct.ufpb.br (A. Subramanian), g.erdogan@bath.ac.uk (G. Erdoğan).

http://dx.doi.org/10.1016/j.cor.2014.10.019
0305-0548/

or postcodes) and allow the development of familiarity, which makes their task easier. In addition, clustered routes have significantly less overlaps. In several cases, the additional routing costs due to cluster constraints are compensated by the ease of implementation and the enhanced driver familiarity.

The literature on the CluVRP is quite limited as of the time of this writing. Sörensen et al. [19] and Sevaux and Sörensen [17] presented an integer programming formulation capable of finding the best Hamiltonian path between each pair of vertices in each cluster. Barthélemy et al. [2] suggested to adapt CVRP algorithms to the CluVRP by including a large positive term M to the cost of the edges between clusters and a cluster and the depot. The CluVRP is solved as a CVRP by means of the algorithm of Clarke and Wright [6] followed by 2-OPT moves and Simulated Annealing (SA). The authors also suggested to dynamically set the penalty M, but observed that the M term interferes with the Boltzmann acceptance criterion of the SA and leads to erratic performance. Computational results were not reported in this initial paper.

Pop et al. [14] described the directed CluVRP as an extension of the *Generalized Vehicle Routing Problem* (GVRP, [7]). The authors adapted two polynomial-sized formulations for the GVRP to the directed CluVRP, but again no computational results were reported. Recently, Battarra et al. [3] proposed exact algorithms for the CluVRP and provided a set of benchmark instances with up to 481 vertices. The best performing algorithm relies on a preprocessing scheme, in which the best Hamiltonian path is precomputed for each pair of endpoints in each cluster. This allows for selecting a pair of endpoints in each cluster rather than the whole path, relegating some of the problem complexity in the preprocessing scheme. The resulting minimum cost Hamiltonian path problems are reduced to instances of the Traveling Salesman Problem (TSP) and optimally solved with Concorde [1]. CluVRP instances of much larger size than the corresponding CVRP instances were optimally solved, thus highlighting the advantage of acknowledging the presence of clusters.

In this paper, we introduce hybrid adaptations of state-of-the-art CVRP metaheuristics for the CluVRP. Rather than rediscovering well-known metaheuristic concepts, we exploit the current knowledge on Iterated Local Search and Hybrid Genetic Algorithms [20,28] and focus our attention on developing efficient problem-tailored neighborhood searches and effectively embedding them into these metaheuristic frameworks. The proposed neighborhood searches aim at (1) better exploiting clustering constraints by means of pruning techniques, (2) exploring larger neighborhoods by means of dynamic programming, (3) reducing the computational time by means of re-optimization, bi-directional search, and data structures. Finally, these experiments lead to further insights on which type of metaheuristic to use for different instance sizes and cluster characteristics.

The remainder of the paper is organized as follows. Section 2 introduces the challenges related to the CluVRP. Sections 3 and 4 describe the proposed metaheuristics and efficient neighborhood-search strategies, whereas Section 5 discusses our computational results. Conclusions are drawn in Section 6, and further avenues of research are discussed.

2. Motivation

Battarra et al. [3] showed that exact algorithms are capable of solving relatively large CluVRP instances. However, the CPU times remain prohibitively long for large-scale or real time applications. In this paper, we exploit the properties of the CluVRP to develop specialized hybrid metaheuristics that take advantage of cluster constraints. Solution quality is assessed by a comparison with exact solutions whenever possible, and among metaheuristics when it is not.

Two recent and successful metaheuristic frameworks are used in this work. The Iterated Local Search (ILS) algorithm of Subramanian [20] is simple and flexible, combining the intensification strength of Local Search (LS) operators and effective diversification through perturbation operators. It proved to be remarkably efficient for many variants of the Vehicle Routing Problem (VRP), including the VRP with Simultaneous Pickup and Delivery [23], the Heterogeneous VRP [12], the Minimum Latency Problem [18], and the TSP with Mixed Pickup and Delivery [21]. The success of ILS is due to an intelligent design of intensification and diversification neighborhoods, as well as a random exploration of the neighborhoods. This latter component allows for extra diversity, and leads to high quality solutions, even when applied to other problems such as scheduling [22].

ILS explores only feasible solutions, and allows for testing the M approach suggested by Barthélemy et al. [2] without possible interferences between M and penalties applied to infeasible solutions. As mentioned in the Introduction, the M approach consists of including a large positive term to all those edges that are connecting clusters and connecting the depot to the clusters. Any CVRP algorithm in which the M is chosen to be large enough returns a CluVRP solution in which the number of penalized edges is minimized, therefore a solution in which the cluster constraint is satisfied. Note that the number of edges connecting clusters or connecting the depot to a cluster is $m+N$ and their penalization can be easily deducted from the solution cost.

One drawback of this transformation is that most VRP neighborhoods consider moves of one or two vertices. These neighborhoods can often not relocate complete clusters, and thus many moves appear largely deteriorating due to M penalties, significantly inhibiting the progress towards higher quality solutions. As shown in this paper, ILS can partly overcome this issue by means of perturbation moves. However, as demonstrated by our computational results, a more clever application of the framework specific to the CluVRP considering relocating and exchanges of whole clusters and intra-cluster improvements produces solutions of comparable quality in considerably less CPU time. In the next section, we describe the ILS and these hybrid algorithms in more details.

The Unified Hybrid Genetic Search (UHGS) currently obtains the best known solutions for more than 30 variants of the CVRP and represents the state-of-the-art among hybrid metaheuristics for VRPs. More precisely, the algorithm successfully solves problems with diverse attributes, such as multiple depots and periods [25], time windows and vehicle-site dependencies [26], hours-of-service-regulations for various countries [8], soft, multiple, and general time windows, backhauls, asymmetric, cumulative and load-dependent costs, simultaneous pickup and delivery, fleet mix, time dependency and service site choice [28], and prize-collecting problems [30], among others. It has been recently demonstrated that several combinatorial decisions, such as customer selections or depot placement, can be relegated directly at the level of cost and route evaluations, allowing to always rely on the same metaheuristic and local search framework while exploring large neighborhoods in polynomial or pseudo-polynomial time [29,30].

Our UHGS implementation is based on the assumption that the costs of the optimal Hamiltonian paths among vertices in the same cluster can be efficiently precomputed as in Battarra et al. [3]. Once these paths and their costs are known, an effective route representation as an ordered sequence of clusters can be adopted, and a