Decomposition and Set Covering Strategies for Large-Scale Heterogeneous Vehicle Routing Problems

Christoph Kerscher[1], Stefan Minner[1,2] Fabien Lehuédé[3], Guillaume Massonnet[3]

[1]Logistics and Supply Chain Management, Technical University of Munich, Munich, Germany
[2]Munich Data Science Institute (MDSI), Technical University of Munich, Munich, Germany
[3]IMT Atlantique, LS2N, UMR CNRS, Nantes, France

February 11, 2025

Keywords: heterogeneous vehicle routing, matheuristic, set covering, decomposition

1 INTRODUCTION

Motivated by the increased demand for parcel shipments in last-mile delivery, a new research stream has emerged in vehicle routing: finding high-quality solutions for very large-scale vehicle routing problems (VRPs) in reasonable time (e.g., Accorsi & Vigo, 2024, Arnold *et al.*, 2019, Queiroga *et al.*, 2021, Santini *et al.*, 2023). While these approaches show the efficiency of applying decomposition and pruning methods in state-of-the-art metaheuristics like FILO by Accorsi & Vigo (2021) and the hybrid genetic search by Vidal *et al.* (2012), they focus on the capacitated VRP (CVRP) projected in the Euclidean plane. Practitioners, however, operate under more complex conditions, including the use of heterogeneous vehicle fleets, navigating through nested road networks, and the necessity to meet customers' delivery time windows.

Furthermore, real-time planning typically requires solutions within minutes rather than hours. However, given the maturity of research on CVRP instances, recent improvements tend to be marginal, often resulting in solution quality enhancements of less than one percent, even after several hours of computation (e.g., Queiroga *et al.*, 2021, Santini *et al.*, 2023). While high-quality route plans are essential, extending computation times to achieve slight improvements is often impractical, especially when such minor gains are likely to be negated as they are often subject to change due to unforeseen events during route execution.

The goal of this study is to address these issues by developing a method that finds high-quality solutions fast, i.e., in minutes, to large-scale instances of VRP variants that include a heterogeneous fleet and delivery time windows (HVRPTW). We develop an iterative matheuristic that combines decomposition, route generation, and set covering strategies. Customer- and route-based decomposition divides an instance into multiple smaller subproblems. Those are solved separately by an existing metaheuristic to construct a pool of routes. Then, solving the set covering problem (SCP) yields the best set of routes.

The results of the numerical study indicate that our approach effectively leverages the route generation capabilities of state-of-the-art metaheuristics. By integrating these solvers into a decomposition and set covering framework, we successfully solve large-scale instances of various HVRP variants in minutes. Thus, with this research, we intend to make the following contributions:

1. We develop a matheuristic framework that combines complexity reduction techniques, metaheuristics, and a set covering model for solving large-scale HVRPTW instances in a short amount of time.

2. We perform a numerical study on large-scale variants of the HVRP and report many new best-known solutions for instances introduced by Bräysy *et al.* (2009) and Pessoa *et al.* (2018).

3. We conduct a case study on HVRPTW instances using real-world road fleet composition rules to evaluate the performance of the matheuristic in practical routing scenarios.

2 PROBLEM STATEMENT

Following Koç *et al.* (2016), we focus on the two most prominent versions of the HVRP: the fleet size and mix VRP (FSMVRP) and the heterogeneous fixed fleet VRP (HFVRP). In both versions, a set of customers with known demands must be served by a set of vehicles, i.e., a fleet, which is located at a depot. The FSMVRP assumes that the number of available vehicles of each type is equal to the number of stops to be served. In contrast, the HFVRP operates with a predetermined and, therefore, limited number of vehicles per type. Both FSMVRP and HFVRP can be further categorized into variants where either the fixed costs, variable costs, or both are vehicle type-dependent. In all variants, the capacity of the vehicles is limited and type-dependent. The objective in each case is to minimize operational transportation costs, which are the sum of the fixed and variable costs of the vehicles used.

A solution is considered feasible if each customer stop is visited exactly once by a single vehicle. If a delivery time window is specified, the vehicle must begin service within that period, even if it requires waiting due to early arrival. The time window of the depot defines the operational period within which all vehicles must start and complete their routes. Additionally, the number of routes assigned to a particular vehicle type cannot exceed the number of available vehicles of that type, and the total demand assigned to a single vehicle must not exceed its capacity.

3 SOLUTION APPROACH

Our matheuristic consists of the following four recurring steps.

1. **Decomposition:** Every iteration starts with decomposing an instance using customer- and route-based fuzzy c-medoids clustering (Bezdek, 1981) for varying values of c, representing the number of clusters. In fuzzy clustering, data objects are not strictly assigned to a single cluster. Instead, a fuzzy value indicates the relative degree of membership between an object and a cluster based on its similarity to the cluster medoid. Thus, by setting an appropriate threshold, customers can be assigned to multiple clusters to create overlapping subproblems. This approach allows us to generate additional routes in areas where customer stops are densely concentrated and where an intuitive spatial split of them is not applicable. We add route-based decomposition to increase the variety of subproblems. Here, the customers of incumbent routes are combined into multiple subproblems based on the similarity between routes. Once the clusters are formed and the depot vertex is duplicated and added to every cluster, the fleet has to be split among them. We formulate this as a mixed-integer linear program (MILP) that minimizes the dissimilarity of customers of the same cluster. The MILP is constrained by the vehicles' availability and the fulfillment of the clusters' demand.

2. **Routing:** The subproblems are solved independently for each available vehicle type, either as a CVRP or VRPTW, depending on the structure of the original problem. In cases where the capacity of a vehicle type is insufficient to meet the demand of a particular customer, that customer is excluded from the respective subproblem and vehicle type combination. This removal ensures that a feasible set of routes can always be obtained for the remaining

customers. Alternatively, we reformulate the VRPs in the subproblems as prize-collecting problems with a heterogeneous fleet. This approach ensures compatible routes among the different vehicle types, even when solved separately.

3. **Set Covering:** The SCP is formulated as an integer program that minimizes the sum of the total costs of the selected routes while ensuring that (i) each customer is visited at least once and (ii) the number of routes of a particular vehicle type does not exceed its maximum available number. Thus, solving the SCP of the collected routes yields the best route set out of the route pool created in the previous step.

4. **Improving:** A local search step is applied to improve the route set obtained from the SCP. The search includes removing the customers that are served in more than one route. Then, the next iteration starts again at step 1 by decomposing the problem.

The matheuristic terminates when its stopping criterion is met, i.e., the runtime limit is reached or completing a specified number of iterations without improvement.

4 NUMERICAL STUDY

Setup: We apply our matheuristic to the HVRP instances of Pessoa *et al.* (2018) and the FSMVRPTW instances of Bräysy *et al.* (2009). This numerical study examines large-scale routing scenarios, explicitly focusing on instances with a minimum of 300 customers. To simulate realistic route planning conditions, each instance is given a maximum runtime of 600 seconds, regardless of its size. The parameter c, which represents the number of clusters, is defined in such a way that each subproblem is comprised of customers between 50 and 200. In every iteration, the clustering technique in the decomposition step and the formulation of the problem in the routing step are selected uniformly at random. Each experiment is repeated five times with different random seeds. The performance of the proposed approach is evaluated based on the error gap compared to the best-known solutions (BKS) reported in the literature. All experiments are run on a single thread of an Intel(R) Xeon(R) Platinum 8160 CPU 2.1 GHz processor with **2.8**GB of RAM running Ubuntu 20.04 LTS. The matheuristic is implemented in Python 3.9.7. We use PyVRP 0.9.1 by Wouda *et al.* (2024) to generate routes and Gurobi 11.0.3 to solve the SCP.

Preliminary Results: It should be emphasized that the reported results are based on a prototypical version of the matheuristic. Table 1 presents the mean error gap relative to the BKS and the number of instances for which our method improves the existing BKS (column *no. of new BKS*). The results are split by instance size (column *no. of customers*) and VRP variant. For the 57 HVRP instances with at least 300 customers, we identify 11 new BKS. Here, the overall mean error gap across all tested instances is 3.39%, indicating that the matheuristic converges rapidly to high-quality solutions. For the 480 HVRPTW instances by Bräysy *et al.* (2009), we report 31 new best-known solutions and achieve an overall mean error gap of 2.20%. Notably, for both datasets, solution quality does not deteriorate significantly with increasing problem size, demonstrating that the proposed matheuristic is highly scalable and well-suited for larger instances.

Table 1 – *Mean Error Gap to BKS and Number of New BKS Found by Matheuristic*

	FSMVRP & HFVRP		FSMVRPTW	
no. of customers	mean e-gap	no. of new BKS	mean e-gap	no. of new BKS
$\leq$ 400	3.17%	4	2.11%	6
401-600	3.44%	4	2.26%	6
601-800	2.88%	1	2.05%	7
$>$ 800	4.09%	2	2.34%	12
Total	**3.39%**	**11**	**2.20%**	**31**

5 OUTLOOK

Despite the fact that this research project is still ongoing, the presented results demonstrate the competitive performance of the proposed matheuristic, achieving solutions that are close to or improving the current BKS. However, the HVRP dataset by Pessoa *et al.* (2018) is relatively new and has not yet received much attention from the research community. Therefore, it is likely that the solutions can be further improved, and that we will present more advanced results by the time of the conference. As a next step, we extend the runtime to several hours to simulate overnight route planning processes. Thereby, we provide an analysis of the algorithm's complete convergence profile. Furthermore, we analyze the meta-data of the routes of the final solution, e.g., in which iteration they were generated, which clustering technique was applied, and which reformulation was used in the routing phase. This analysis will provide valuable insights into which customers are particularly challenging to incorporate into the overall route plan. Lastly, we apply our matheuristic to HFVRPTW instances based on actual fleet composition scenarios to demonstrate its applicability in real-world delivery use cases.

References

Accorsi, Luca, & Vigo, Daniele. 2021. A fast and scalable heuristic for the solution of large-scale capacitated vehicle routing problems. *Transportation Science*, **55**(4), 832–856.

Accorsi, Luca, & Vigo, Daniele. 2024. Routing one million customers in a handful of minutes. *Computers & Operations Research*, **164**, 106562.

Arnold, Florian, Gendreau, Michel, & Sörensen, Kenneth. 2019. Efficiently solving very large-scale routing problems. *Computers & Operations Research*, **107**, 32–42.

Bezdek, James C. 1981. *Pattern Recognition with Fuzzy Objective Function Algorithms*. Boston: Springer US.

Bräysy, Olli, Porkka, Pasi P, Dullaert, Wout, Repoussis, Panagiotis P, & Tarantilis, Christos D. 2009. A well-scalable metaheuristic for the fleet size and mix vehicle routing problem with time windows. *Expert Systems with Applications*, **36**(4), 8460–8475.

Koç, Çağrı, Bektaş, Tolga, Jabali, Ola, & Laporte, Gilbert. 2016. Thirty years of heterogeneous vehicle routing. *European Journal of Operational Research*, **249**(1), 1–21.

Pessoa, Artur, Sadykov, Ruslan, & Uchoa, Eduardo. 2018. Enhanced branch-cut-and-price algorithm for heterogeneous fleet vehicle routing problems. *European Journal of Operational Research*, **270**(2), 530–543.

Queiroga, Eduardo, Sadykov, Ruslan, & Uchoa, Eduardo. 2021. A POPMUSIC matheuristic for the capacitated vehicle routing problem. *Computers & Operations Research*, **136**, 105475.

Santini, Alberto, Schneider, Michael, Vidal, Thibaut, & Vigo, Daniele. 2023. Decomposition strategies for vehicle routing heuristics. *INFORMS Journal on Computing*, **35**(3), 543–559.

Vidal, Thibaut, Crainic, Teodor Gabriel, Gendreau, Michel, Lahrichi, Nadia, & Rei, Walter. 2012. A hybrid genetic algorithm for multidepot and periodic vehicle routing problems. *Operations Research*, **60**(3), 611–624.

Wouda, Niels A., Lan, Leon, & Kool, Wouter. 2024. PyVRP: A High-Performance VRP Solver Package. *INFORMS Journal on Computing*, **36**(4), 943–955.