

MASTERARBEIT | MASTER'S THESIS

Titel | Title

A Multi-Compartment Vehicle Routing Problem with Loading and Unloading Costs

verfasst von | submitted by

Pt. (T.E.) Vasileios Rafail Karasavvidis

angestrebter akademischer Grad | in partial fulfilment of the requirements for the degree of
Master of Science (MSc)

Wien | Vienna, 2026

Studienkennzahl lt. Studienblatt |
Degree programme code as it appears on the
student record sheet:

UA 066 915

Studienrichtung lt. Studienblatt | Degree
programme as it appears on the student
record sheet:

Masterstudium Betriebswirtschaft

Betreut von | Supervisor:

Ass.-Prof. Dr. Christian Tilk

Abstract

This thesis addresses the multi-compartment vehicle routing problem with loading and unloading costs (MCVRP_LU), described by Hübner and Ostermeier (2019), treating it as a specification to be implemented. Food retailers supply their shops from a distribution centre with goods from various temperature segments. Multi-compartment vehicles allow for the joint delivery of these segments, albeit at the cost of an additional loading cost for each activated compartment, meaning that the number of active compartments simultaneously influences the costs of loading, transport and unloading. The published model is reformulated to explicitly state its assumptions, and the two-stage solution method, a parallel savings construction followed by a large neighbourhood search with Shaw removal, regret-k insertion and Record-to-Record Travel acceptance, is implemented in C++ based solely on the published description and the set of parameters. The solver is applied to the seven instances for which the authors have published both the instance data and the corresponding solution files, taking into account three levels of handling costs and five independent runs in each case, resulting in a total of 105 runs. The structure of the solution and the composition of the costs are reproduced with a high degree of accuracy: The number of routes agrees to within four routes in every instance; the average number of active compartments per route deviates by a maximum of 0.28 in six of the seven instances; and the published cost shares for loading, transport and unloading match, apart from a deviation of 0.9 percentage points.

Abstract

Diese Arbeit behandelt das Multi-Compartment Vehicle Routing Problem mit Be- und Entladekosten (MCVRP_LU), das von Hübner und Ostermeier (2019) beschrieben wurde, und versteht dieses als Spezifikation für eine Implementierung. Lebensmitteleinzelhändler beliefern ihre Filialen von einem Distributionszentrum aus mit Waren verschiedener Temperatursegmente. Mehrkammerfahrzeuge ermöglichen die gemeinsame Auslieferung dieser Segmente, verursachen jedoch für jedes aktivierte Fahrzeugabteil zusätzliche Beladekosten. Die Anzahl der aktivierten Abteile beeinflusst somit gleichzeitig die Kosten für Be- und Entladung sowie den Transport.

Das veröffentlichte Modell wird neu formuliert, um die zugrunde liegenden Annahmen explizit darzustellen. Darüber hinaus wird das zweistufige Lösungsverfahren, bestehend aus einer parallelen Savings-Konstruktion und einer Large-Neighbourhood-Search mit Shaw Removal, Regret-k-Insertion und Record-to-Record-Travel-Akzeptanz, ausschließlich auf Grundlage der veröffentlichten Beschreibung und des angegebenen Parametersatzes in C++ implementiert.

Der entwickelte Solver wird auf den sieben Instanzen angewendet, für die die Autoren sowohl die Instanzdaten als auch die zugehörigen Referenzlösungen veröffentlicht haben. Dabei werden drei Stufen von Umschlagskosten sowie jeweils fünf unabhängige Durchläufe pro Instanz berücksichtigt, sodass insgesamt 105 Läufe durchgeführt werden.

Die Struktur der Lösungen sowie die Zusammensetzung der Kosten werden mit hoher Genauigkeit reproduziert: Die Anzahl der Touren stimmt in jeder Instanz bis auf höchstens vier Touren mit den veröffentlichten Ergebnissen überein. Die durchschnittliche Anzahl aktiver Fahrzeugabteile pro Tour weicht in sechs der sieben Instanzen um höchstens 0,28 ab. Zudem werden die veröffentlichten Kostenanteile für Beladung, Transport und Entladung mit Ausnahme einer Abweichung von 0,9 Prozentpunkten nahezu exakt nachvollzogen.

Statement of use of Artificial Intelligence

Artificial intelligence-based applications were used as tools in the preparation of this thesis. Claude (Anthropic) and Gemini (Google) were used to assist with reviewing specific passages, checking the placement of quotations, analysing the structure and organisation of the argument. In addition, Grammarly and DeepL were used for linguistic and grammatical checking, as well as to improve the overall readability of the text passages.

All the applications mentioned above were used exclusively as a support tool and in consultation with the thesis supervisor. Their use was in accordance with the University of Vienna's current guidelines on artificial intelligence and respected scientific independence: they served to facilitate the linguistic, technical and practical aspects of the work and did not in any way replace the author's own scientific work. All results served solely as technical or linguistic support and were critically reviewed by the author, adapted and placed in context. The formulation of the problem, the design and application of the methodology for its solution, the interpretation of the results, and the final development of the argumentation were carried out independently. The author assumes sole responsibility for all content, code and results of this master's thesis.

Table of Contents

1. Introduction.....	1
1.1 Motivation.....	1
1.2 The multi-compartment vehicle routing problem with loading and unloading costs.....	2
1.3 Research problem.....	2
1.4 Objectives and research questions	4
1.5 Approach and scope	4
1.6 Structure of the thesis.....	5
2. Practical Background.....	5
2.1 Grocery distribution and multi-temperature logistics	6
2.2 Single-compartment versus multi-compartment distribution.....	6
2.3 Decision and relevant cost component.....	7
2.4 Technical configuration of multi-compartment vehicles	9
2.5 Related work	9
3. Model and methodology	11
3.1 Problem Definition.....	12
3.2 Modelling assumptions	14
3.3 Cost components.....	14
3.4 Mathematical formulation.....	16
3.4.1 Complete Model.....	16
3.4.2 Decision variables	16
3.4.3 Objective function.....	17
3.4.4 Constraints	17
3.5 Solution methodology	18
3.5.1 Construction heuristic: parallel savings	18
3.5.2 Shaw removal.....	19
3.5.3 Regret-k insertion.....	21
3.5.4 Record-to-Record Travel acceptance	23
3.5.5 Parameter setting.....	23
4. Results.....	24
4.1 Setup and comparison	25
4.2 Total cost against the published reference value	26
4.3 Solution structure	28
4.4 Cost composition	30
4.5 Runtime.....	30

4.6 Stability across runs	32
5. Conclusion	33
5.1 Summary of the study	33
5.2 Answers to the research questions	34
Appendix.....	35
List of Figures	35
List of Tables.....	35
List of Abbreviations.....	36
Results for standard, 50% and 100%.....	37
References.....	44

1. Introduction

1.1 Motivation

Grocery retailing businesses operate daily on thin commercial margins and distribution is one of the largest cost blocks in that sector. In the grocery supply chain, distribution accounts for about one-fifth of the total logistics cost, while the net retail margin commonly remains under two percent (Hübner & Ostermeier, 2019). For that reason, even the smallest reductions in transportation, loading and unloading translate into substantial absolute savings and thus the design of efficient delivery operations is an issue of direct economic relevance and not a mere academic concern.

Grocery distribution has the peculiarity of simultaneously dealing with multiple temperature regimes. Products are stored and transported in different zones, e.g. deep frozen, chilled and ambient, and the permissible temperature ranges are subject to legislative requirements (Hübner & Ostermeier, 2019). In the European Union, mandatory transport temperatures range from -20°C to -18°C for deep frozen products to $+2^{\circ}\text{C}$ to $+7^{\circ}\text{C}$ for chilled goods (Frank et al., 2021). Historically, this segmentation meant that retailers would need to use a different vehicle and thus a different tour to serve each segment since a single conventional truck could only serve one temperature regime at a time.

Multi-compartment vehicles (MCVs) are changing the picture. The loading area is divided into temperature-specific compartments so that a single MCV can transport several product segments in one tour and deliver them jointly to a point of sale (Hübner & Ostermeier, 2019). The consolidation achieved is especially interesting for the urban deliveries and for the points of sale with difficult access, where the dispatch of several single-compartment vehicles (SCVs) is extremely expensive and operationally inconvenient. At the same time, MCVs create a new tension in planning. If more than one shipping gate is approached at the distribution centres (DC), because several segments are collected, the loading effort increases. If segments are combined for one customer, the number of stops and the unloading effort are reduced. The number of compartments activated on a tour therefore drives loading, transportation and unloading costs simultaneously, and these effects pull in opposite directions.

1.2 The multi-compartment vehicle routing problem with loading and unloading costs

The routing question which arises in that setting is one variant of the classical capacitated vehicle routing problem (CVRP) and is referred to as a multi-compartment vehicle routing problem (MCVRP). In the formulation of Hübner and Ostermeier (2019), the compartments are completely flexible, which means that determined per vehicle and tour, subject only to the total vehicle capacity. The central modelling contribution of that work is the explicit integration of compartment-dependent loading and unloading costs into the objective, which yields the model the authors denote MCVRP_LU. In this model, loading costs increase with each additional compartment, as each segment must be collected from a different door, while unloading costs per stop are lower when several segments are delivered at the same time. Given that the cost factors depend on the number of active compartments, which, in turn, is a separate decision, the objective function is non-linear (Hübner & Ostermeier, 2019).

To solve practically sized instances of up to roughly four hundred customer deliveries, Hübner and Ostermeier (2019) propose a two-stage heuristic. An initial feasible solution is produced with the parallel savings construction of Clarke and Wright (1964) and this solution is then improved by a large neighbourhood search (LNS) that repeatedly removes and reinserts orders. The elimination phase follows Shaw's (1997) similarity-based logic, whilst the reintegration phase employs the regret operator criterion modelled on that of Ropke and Pisinger (2006) and the selection of candidate solutions is based on a 'Record-to-Record Travel' (RRT) criterion (Dueck, 1993), supplemented by a restart mechanism that diversifies the search when progress stalls. A distinctive feature of the approach is that the neighbourhood is defined over the movement of orders between vehicles rather than over the movement of customers, which reflects the fact that a single customer may receive several segment-specific orders on different tours.

1.3 Research problem

The published article provides a mathematical formulation, a pseudocode description of the LNS and its operators, a documented set of algorithm parameters and a collection of test examples, all of which are publicly available along with the reference solutions. Taken together, these elements constitute a specification of a solution method. However, it is not immediately

obvious whether this specification alone is sufficient to reproduce the described behaviour of the solver, as various decisions regarding modelling and implementation are described only in general terms. This thesis takes this issue as its starting point. It treats the article by Hübner and Ostermeier (2019) as a specification to be implemented, rather than as a result to be referenced.

Comparing studies in this field depends on common benchmark instances with known solutions, making the collection and cataloging of such datasets a crucial focus in vehicle routing research (Gunawan et al., 2021). However, having a published set of instances is only part of what's needed for comparison. Equally important is the ability to reconstruct the method that generated the reported solutions. Kohan et al. (2026) observe that recent studies on multi-compartment routing with flexible compartments rarely share underlying implementations, which hinders reproducibility and makes meaningful comparison difficult. When results cannot be reproduced from the description alone, verifying, comparing or building upon those methods becomes significantly more difficult. The MCVRP_LU model is suitable as the subject of this study for two reasons. First, it is clearly defined, so that the reconstruction can be evaluated based on essential structural expectations and not solely based on aggregate data; second, the authors have made their examples and reference values publicly available, which allows for a comparison under identical conditions.

The difficulty of verifying published results of this kind is not a problem limited to the present case. Kohan et al. (2026) point out that none of the more recent papers on multi-compartment routing with flexible compartments have published their implementations as open source, which, in their view, limits both reproducibility and comparative evaluation; they state that they are unable to compare their results directly with those of Ostermeier and Hübner (2018), as neither the implementation, nor the test cases, nor the individual results are available. Hübner and Ostermeier (2019) are considerably more generous in this regard, as both the test cases and the corresponding solution files have been published and it is precisely this transparency that enables the present comparison. What has not been published is the implementation itself, as well as the intermediate values that would document how an execution arrives at its result. Consequently, any replication must reconstruct the solution method solely on the basis of the text of the article and any decision regarding the model or the algorithm that the text leaves open must be made by the person carrying out the reconstruction, rather than being derived from the source.

1.4 Objectives and research questions

The overall objective of this thesis is to determine to what extent the MCVRP_LU model and its solution method can be reconstructed solely based on the published description, as well as to identify the points where that description requires interpretation. This objective is addressed through two research questions.

First: Can a solver, which has been constructed in accordance with the specifications of Hübner and Ostermeier (2019) and is run on the examples they have published, produce results that agree with the reported reference values, both in terms of the numerical target values achieved and in terms of the structural characteristics of the solutions, such as the number of routes, the number of stops, and the average number of compartments used?

Second: In what respects does the published description implicitly allow for specific modelling decisions or algorithmic choices, and to what extent do the logical outcomes of these decisions influence the reconstructed results?

1.5 Approach and scope

From a methodological perspective, the thesis is divided into two interconnected phases that reflect the structure of the original paper. The first phase includes the reconstruction of the model, in which the mathematical formulation is reformulated and the modelling assumptions that remained implicit in the article are explicitly stated, so that the solver is based on a clear foundation. The second phase consists of the algorithmic and computational reconstruction, during which parallel savings construction and the LNS, with the Shaw removal operator, its reintroduction based on a regret-k insertion, the reset mechanism and RRT acceptance, are implemented and subsequently applied to the cases. Each case is solved iteratively according to the protocol of repeated runs from the original study and the results are compared with the reference values.

The scope of the study is based on the data published by Hübner and Ostermeier (2019). The main comparison set includes the case families for which the authors publish both the cases and the corresponding reference solutions: Test A.2, in which the average order size varies between three, five, seven, and nine transport units (TU) across 199 customer locations and 400 orders, Test B, which has the same scope but shorter distances between clusters and Tests C.1 and C.2, in which the scope is reduced to 106 locations with 198 and 106 orders,

respectively. All main cases use three product segments and a vehicle capacity of 33 TU, and are solved, respectively, at the standard level, as well as at 50% and 100% increased handling costs (loading and unloading), while for each case and level, five independent runs are performed with different seeds.

1.6 Structure of the thesis

The remainder of the thesis is divided into four chapters. Chapter 2 describes the practical context, namely the structure of distribution networks in the food sector, the difference between single-compartment and multi-compartment vehicles, the cost components that must be taken into account in route planning, and the technical configuration of the vehicles and reviews the related literature. Chapter 3 revisits the MCVRP_LU problem formulated by Hübner and Ostermeier (2019) and discusses the problem definition, the model assumptions, the cost components and the mathematical formulation; it then describes the solution method in the order in which its components are combined, namely the construction of parallel savings, Shaw's elimination, the insertion of 'regret-k', the 'Record-to-Record Travel' acceptance rule and the set of parameters. Chapter 4 reports on the reconstruction. It describes the instances, cost levels and comparison protocol, and then compares the reconstructed solutions with published benchmark solutions in terms of total costs, solution structure, cost composition, running time and stability across different iterations. Chapter 5 summarises the study, answers the two research questions.

2. Practical Background

This chapter explains the operational context in which the routing problem arises. It begins by outlining the structure of food distribution networks and the role played by multi-compartment vehicles within them; it then compares operations using a single compartment with those using multiple compartments and finally, it describes the cost components that distinguish this problem from a standard routing problem, together with the technical flexibility of the vehicles. The description is deliberately kept informal and serves as a basis for the formal model in Chapter 3. Section 2.5 then positions the problem within the existing literature and states the gap this thesis addresses.

2.1 Grocery distribution and multi-temperature logistics

The majority of products sold by grocery retailers are shipped to stores through distribution centres (DCs) rather than a direct supplier-to-store transportation system (Ferne & Sparks, 2009, cited in Hübner & Ostermeier, 2019; Kuhn & Sternbeck, 2013, cited in Hübner & Ostermeier, 2019). European discount chains and supermarkets operate their own regional networks of DCs, each of which ships the full range of products to a designated set of stores (Hübner & Ostermeier, 2019). Typically, one DC serves between 50 and 400 outlets (Glatzel, Großpietsch, and Hübner, 2012, cited in Hübner & Ostermeier, 2019). This range influences the scale of routing challenges faced in practice and explains why the original study focused on problem sizes of up to 400 customers.

A key feature of this environment is temperature segregation. Products are stored, picked and delivered from temperature-controlled zones within the distribution centres (DC) due to regulatory requirements that specify permissible transportation temperatures, which vary for deep-frozen, chilled and ambient products (Hübner & Ostermeier, 2019). Consequently, the DCs are organised by temperature zones, with products from each zone loaded at specific gates. This structural organisation, however, leads to a conflict: serving multiple temperature zones necessitates managing multiple gates.


2.2 Single-compartment versus multi-compartment distribution

The traditional model for distribution, which uses single-compartment vehicles (SCVs), states that a vehicle starts its journey from only one shipping gate per temperature. Once the vehicle loads all the goods in one specific segment, it delivers to all the customers who are involved in the same tour. If a customer has ordered more than one product segment, but the other segment has a different temperature requirement, a second tour will be conducted with another vehicle, since SCVs can accommodate only one temperature at a time.

An MCV does not impose this constraint. It visits multiple shipping gates at the DC and sorts shipments into separate compartments according to their temperature characteristics, making it possible to deliver several segments to the outlets using just one round trip (Hübner & Ostermeier, 2019). The viability of this strategy has been demonstrated in both practice and research the most common example mentioned in research, is that of German logistics operator Lekkerland, which received acclaim in 2013 for implementing multi-temperature logistics and

the bundling effect in its operations (Grünrock-Kern, 2013, cited in Hübner & Ostermeier, 2019). Retailers have seen the increasing benefits of investing in an MCV, especially in cases when deliveries are made to urban locations and to outlets with restricted access, where sending one truck for each segment is either costly or operationally difficult (Hübner & Ostermeier, 2019). Ostermeier et al. (2021) have quantified the situation: while three temperature zones require three separate routes, each with one compartment and a total of seven stops, a single route with several compartments and four stops can meet the same demand, provided there is sufficient capacity.

Figure 1: Load structure and delivery route of a multi-compartment vehicle (adapted from Hübner & Ostermeier, 2019)



2.3 Decision and relevant cost component

The switch from SCV to MCV does not simply change which vehicle is employed; it changes how the operations are costed. Hübner and Ostermeier (2019) analyse the pertinent processes by looking at a time and motion study performed in conjunction with a retailer and distinguish two process effects and, alongside these, the transportation cost.

First, the cost-increasing effect: when more than one segment is carried, the vehicle must make multiple stops at the DC to deliver each segment to the relevant gate; this adds both travel and setup time during loading. Loading setup time consists of travel to the relevant gate and the time required to open and close both the compartment and its partition. Loading cost can then best be considered as setup cost, which increases with the number of compartments, since the number of compartments directly reflects the number of gate setups performed at the DC (Hübner & Ostermeier, 2019).

Second, the cost-decreasing effect arises because MCVs can deliver multiple segments at a single customer stop, thereby reducing the number of stops and unloading setups at delivery points (Hübner & Ostermeier, 2019). Unloading costs can themselves be separated into two parts: the generic, per-stop costs like manoeuvring to the ramp or opening the truck, which both SCVs and MCVs share and thus depend only on the number of stops, that is, on the number of customers visited. These general per-stop costs account for the bulk of decision-relevant unloading effort. Compartment-specific costs reflect the opening and closing of compartments and are minimal, requiring only a few seconds per stop and being almost independent of whether one or more compartments are opened once more than one is in use. The model therefore needs only to distinguish between MCVs and SCVs, not to differentiate unloading by compartment or by customer (Hübner & Ostermeier, 2019).

Third, transportation cost. The greater routing flexibility of the MCV can translate into a reduced total driving distance compared with using an SCV when multiple segments can be delivered on a single route (Hübner & Ostermeier, 2019).

These factors collectively create the main trade-off in the problem. Adding another segment increases compartment load and loading costs but can lower unloading costs via delivery consolidation and may reduce overall transport expenses through network simplification. The best balance is not predetermined and must be identified during routing, especially by deciding how many compartments to load in each tour. Two points are important to bear in mind: firstly, the physical moving of goods, e.g. roll-cages, is required independently of the number of compartments and secondly, load and unload costs together represent a significant portion of all operating costs, up to one fifth, that cannot be ignored when routing vehicles (Hübner & Ostermeier, 2019).

2.4 Technical configuration of multi-compartment vehicles

For the MCV, the compartments are completely flexible; neither their size nor the total number is specified before the tour planning process. The compartments can be deactivated completely by removing the compartment wall. The compartment sizes are adjusted to match loaded quantities exactly. For the tests, the vehicle capacity is set to 33 TUs, so each compartment may hold between 0 and 33 TU, where the sum of all compartment sizes must not exceed the total capacity of the vehicle. Current MCV designs additionally permit rear loading, because the partition runs lengthwise along the loading area rather than across it (Hübner & Ostermeier, 2019).

2.5 Related work

Ostermeier et al. (2021) provide the reference overview for this class of problems and structure it according to three compartment-related dimensions: how flexible the compartment sizes are, whether items are assigned to specific compartments and whether items are compatible with one another. Most studies deal with only one or two of these dimensions. The dimension relevant here is the first. In problems involving fixed cargo compartments, the cargo space is divided up in advance and route planning merely determines which orders are delivered where. In problems involving flexible cargo compartments, both the number of compartments and their size are decision variables, limited only by the vehicle's total capacity. The MCVRP_LU belongs to the second group.

Derigs et al. (2011) provide the first complete formulation of the MCVRP problem with flexible compartments and present a solver suite that combines construction and improvement heuristics and large neighbourhood search. However, their focus is exclusively on transport distance. This is typical of earlier work: the number of compartments determines which orders a vehicle can share, but does not incur any costs of its own, meaning that a vehicle with three active compartments is treated in exactly the same way as one with just a single compartment. Hübner and Ostermeier (2019) close this gap. Based on a time and motion study at a food retailer, they show that activating a compartment involves a journey to the gate and a set-up phase at the distribution centre, whilst the delivery of multiple segments at a single stop saves a separate journey to the point of sale. Their model takes both effects into account, so that the

number of active compartments becomes a decision with its own costs and is no longer merely a by-product of route planning. This is the contribution that the present thesis reconstructs.

The model has since been extended in several directions. Ostermeier and Hübner (2018) incorporate vehicle selection, allowing the planner to choose between single-compartment and multi-compartment vehicles, rather than assuming a homogeneous fleet. Frank et al. (2021) integrate route planning for multi-compartment vehicles into the joint planning of weekly delivery patterns, so that the decision on which segments are transported together is made across multiple time periods. Ostermeier (2024) shifts the scope of application to convenience stores, adds multiple time slots, service times, driver breaks and maximum tour durations and uses the target values from Hübner and Ostermeier (2019) as a benchmark. More recently, Masmoudi et al. (2026) generalise flexible cargo compartments to reconfigurable vehicles with modular, dynamically adaptable cargo compartments, whilst Kohan et al. (2026) address the tactical counterpart, namely the size and composition of a heterogeneous fleet with multiple cargo compartments. Outside the food retail context, Lv and Xu (2023) formulate a multi-objective, time-dependent variant with time windows and define, a point particularly worth emphasising here, the loading costs as proportional to the number of active compartments and the unloading costs as a fixed amount per stop, which depends solely on the vehicle type. This corresponds to the cost structure reconstructed in this study.

The heuristic itself was not developed from scratch, but was assembled from tried-and-tested components. The initial solution is derived from the savings algorithm by Clarke and Wright (1964). A large neighbourhood search, whose removal operator follows Shaw's (1997) similarity logic and whose insertion operator uses the regret criterion of Ropke and Pisinger (2006), with acceptance governed by Dueck's (1993) 'Record-to-Record Travel' rule. These components recur frequently in the literature on multi-compartment models. Chen et al. (2019) apply modified Shaw removal to a variant for the cold chain, and Eshtehadi, Demir and Huang (2020) develop an adaptive large-neighbourhood search for urban logistics using multi-compartment models. This facilitates reconstruction, as each operator has its own published specification and its behaviour can be verified against the wider literature, rather than being derived from a single source.

Comparisons between different studies are based on shared instances with a known reference solution and the collection of such datasets has itself now become a topic of research in the field of vehicle route planning (Gunawan et al., 2021). The literature on multi-compartment

models only partially meets this requirement at the instance level and hardly at all at the implementation level. Kohan et al. (2026) report that they are unable to compare their results directly with those of Ostermeier and Hübner (2018), as neither the implementation nor the test instances nor the individual results are available. Runtime presents a related difficulty: Coelho, Cordeau and Laporte (2012) demonstrate for a comparable metaheuristic that a large proportion of the computation time is accounted for by a subproblem that is solved repeatedly, meaning that runtime values between implementations are only comparable if both specify this breakdown.

This raises two points: the MCVRP_LU remains an actively used benchmark and its objective values continue to serve as reference points in current research; it is therefore not yet conclusively clear how reliable these values are. Although the model is documented at the level of its formulation, pseudocode, parameters and published instances, documentation at the implementation level is lacking and no independent reconstruction has been reported to date. This work directly addresses this gap: it reconstructs the model and its solution approach solely on the basis of the published description and compares the result with the published reference solutions across multiple dimensions, rather than relying on a single overall figure.

3. Model and methodology

This chapter revisits the multi-compartment vehicle routing problem with loading and unloading costs (MCVRP_LU), as formulated by Hübner and Ostermeier (2019) and outlines the solution approach. Sections 3.1 to 3.4 reiterate the formulation from Hübner and Ostermeier (2019, pp. 287–288). The notation, variable names and equation numbers are exactly retained from the original to maintain traceability of each solver element to its source. Any differences are solely to clarify assumptions that were implicit initially; such changes are clearly marked. Section 3.5 then explains the solution method, outlined according to component assembly: the parallel savings construction (3.5.1), Shaw removal (3.5.2), regret-k insertion (3.5.3), the Record-to-Record Travel acceptance rule with its control mechanism (3.5.4) and the parameter set (3.5.5).


3.1 Problem Definition

The problem is modelled on an undirected, weighted graph $G = (L, E)$. The node set L includes the distribution centre (DC), labelled as 0, and the locations of n customers, with $L_{\text{cust}} = \{1, \dots, n\}$ representing these customer nodes. The set of edges $E = \{(i, j) : i, j \in L, i < j\}$ indicates the connections between locations, each with a non-negative distance cost, cost_{ij} . Since distances are metric, detours through a third location do not reduce the distance compared to direct links. All vehicle routes are closed loops starting and ending at node 0 (Hübner & Ostermeier, 2019).

Demand is expressed more in terms of orders than in terms of total customer demand. The set of orders is denoted by O and the set of product segments by P , where a product segment comprises products that have the same temperature range and which, therefore, can be transported together in a single loading space. Each order $o \in O$ is characterised by three attributes: the customer who placed it, $\text{customer}(o) \in L_{\text{cust}}$; the requested product segment, $\text{product}(o) \in P$; and a strictly positive quantity. The orders of a specific customer j are grouped into the set $\text{ordCust}(j) = \{o \in O : \text{customer}(o) = j\}$, whilst the orders belonging to a specific product segment p are grouped into O_p . Each customer places at least one order and each product segment is requested in at least one order, but it is not necessary for a customer to place orders for all product segments. The same product segment from the same customer cannot be allocated to more than one vehicle (Hübner & Ostermeier, 2019).

This order-based approach is the key feature that distinguishes the MCVRP_LU problem from the classical route-planning problem for vehicles with limited capacity and which determines the design of the search operators in sections 3.5.2 and 3.5.3. Given that a customer may place several orders and that these orders may travel on different vehicles, it is possible that the same customer may be visited several times throughout the solution and the individual unit that is relocated by the search is the order, not the customer (Hübner & Ostermeier, 2019).

Figure 2: Assignment of segment-specific orders to two vehicle tours (adapted from Hübner & Ostermeier, 2019)



Transport is carried out by a homogeneous fleet V . Each vehicle offers the same total capacity $\text{vehCap}_a > 0$ and the same set of compartments C , and the fleet is treated as unbounded, so that vehicle availability never becomes a binding constraint. The total capacity of a vehicle can be subdivided into up to $\bar{c} \leq |C|$ compartments. In the simplest case, the entire capacity constitutes a single compartment; if more compartments are activated, the size of each follows from the total quantity assigned to it. The size of the compartments is therefore flexible and ranges from zero, meaning that the compartment is not used, to the vehicle's total capacity, and since the sizes are adjusted precisely to the quantity being loaded, the use of compartments does not result in any loss of useful capacity. Two product segments cannot share the same compartment (Hübner & Ostermeier, 2019).

A feasible plan, therefore, consists of m closed routes, each assigned to a vehicle. Route t_k is represented as $(0, \text{stop}_1, \dots, \text{stop}_l, 0)$: starting from the distribution centre (DC), it visits an ordered sequence of l customer stops where $\text{stop}_j \in L_{\text{cust}}$ for $j \in \{1, \dots, l\}$, and then returns to the DC (Hübner & Ostermeier, 2019).

3.2 Modelling assumptions

The following assumptions are taken from Hübner and Ostermeier (2019) and are set out here so that the model and its representation are based on a common framework.

Firstly, the fleet is homogeneous and unlimited, that is to say, there are sufficient vehicles available, all with identical capacity, to meet total demand; therefore, vehicle availability is not a determining constraint. Secondly, each order is allocated in its entirety to a single compartment of an individual vehicle, ruling out split deliveries. Thirdly, the dimensions of the compartments, within the vehicle's capacity limits, are continuous; therefore, the only capacity constraint is that the sum of the loaded quantities must not exceed $vehCapa$, while the binding constraint regarding the compartments is the upper limit $\bar{c}$ on the number of compartments. Fourthly, it is not permitted to mix different product segments within the same compartment, so the number of different segments loaded onto a vehicle is equal to the number of active compartments. Fifthly, the distances are symmetric and satisfy the triangle inequality. Sixthly, each route starts and ends at the distribution centre (DC) and the DC itself is not counted as a stop for the purposes of unloading costs. These assumptions are consistent with the published model and are not relaxed in this study.

3.3 Cost components

The objective combines three cost components, each of which is based on the functional analysis summarised in Chapter 2 and standardised by Hübner and Ostermeier (2019).

The first component is the loading cost. Picking up different product segments requires the vehicle to be positioned at different loading points within the distribution centre (DC); therefore, loading is considered a set-up cost that varies depending on the number of active compartments. It is denoted by $lc_{\{d_v\}}$, where d_v is the number of active compartments in vehicle v . The loading cost is incurred once per vehicle upon departure and increases with each additional compartment (Hübner & Ostermeier, 2019). The second component is the transport cost, which is derived from the distance costs $cost_{ij}$ accumulated along the arcs actually travelled by each vehicle during its journey, including the journey from and back to the distribution centre.

The third component consists of unloading costs, which are regarded as set-up costs for each stop; these are denoted by $ulc_{\{d_v\}}$ and are multiplied by the number of customer stops along

the route. As explained in Chapter 2, the proportion of the unloading effort attributable to a specific compartment is negligible and when more than one compartment is used, it is practically constant. Therefore, when it comes to unloading, the model distinguishes only between single-compartment vehicles (SCVs) and multi-compartment vehicles (MCVs), that is, between $d_v = 1$ and $d_v \geq 2$, and makes no distinction between individual compartments (Hübner & Ostermeier, 2019).

Subsequent studies have also arrived at a similar interpretation of unloading costs. Lv and Xu (2023) build on the earlier vehicle selection model by Ostermeier and Hübner (2018) and formulate loading costs as proportional to the number of active compartments, whilst unloading costs are defined as a fixed amount per stop, depending solely on the vehicle type, with no other dependencies on the number of compartments involved. This corresponds to the distinction made here between SCV and MCV. However, the overlap is limited, as Lv and Xu (2023) draw on the 2018 formulation of vehicle selection rather than directly on Hübner and Ostermeier (2019) and do not indicate that they have attempted to replicate the study; this should be interpreted as meaning that an independent author has arrived at a related cost structure and not as a confirmation of it.

The numerical values of these components are not published as absolute monetary amounts, as they are subject to a confidentiality agreement with the distributor participating in the study; for this reason, Hübner and Ostermeier (2019) express them as ratios relative to the cost of transporting a unit of distance. In the standard cost model, the loading cost of a vehicle depends on the number of active compartments and is reported by Hübner and Ostermeier (2019) as a multiple of the single compartment value of 2.60, namely 2.60 for one active compartment, $2.06 \cdot 2.60$ for two active compartments, and $3.06 \cdot 2.60$ for three active compartments, which corresponds to loading a further compartment adding roughly one single compartment setup. The unloading cost per stop is 2.05 for a single compartment vehicle and $2.05 \cdot 1.12$ for a vehicle with two or more active compartments, where the factor of 1.12 reflects the 112% unloading cost figure reported for multi-compartment vehicles. To analyse the sensitivity of routing decisions with regard to the relative weight of freight handling compared to the route, two additional adjustments increase loading and unloading costs by 50% and 100%, respectively, while transport costs remain unchanged; these scenarios are termed ‘standard’, ‘+50%’ and ‘+100%’ (Hübner & Ostermeier, 2019). The implications of representing these parameters solely as ratios are revisited in Section 4.2 and in Section 5.2.

3.4 Mathematical formulation

The model seeks to minimise all costs such as loading, transportation and unloading by concurrently optimising the number of compartments per vehicle, allocating segment-specific orders to these compartments, and sequencing the routes. Equations (1) to (14), from Hübner and Ostermeier (2019, pp. 287–288), are included with their original numbering, enabling straightforward referencing.

3.4.1 Complete Model

The objective function (Hübner & Ostermeier, 2019) minimises the total cost across all vehicles:

$$= \sum_{v \in V} \left[lc_{d_v} + \sum_{i \in L} \sum_{j \in L} cost_{ij} \cdot b_{ijv} + ulc_{d_v} \left(\sum_{i \in L} \sum_{j \in L_{cust}} b_{ijv} \right) \right] \quad (1)$$

Subject to:

$$\sum_{j \in L_{cust}} b_{0jv} \leq 1 \quad v \in V, \quad (2)$$

$$\sum_{i \in L} b_{ihv} = \sum_{j \in L} b_{h jv} \quad v \in V, h \in L, \quad (3)$$

$$u_{iv} - u_{jv} + |L| \cdot b_{ijv} \leq |L_{cust}| \quad v \in V, i \in L, j \in L_{cust}, \quad (4)$$

$$u_{0v} = 1 \quad v \in V, \quad (5)$$

$$\sum_{o \in O} \sum_{c \in C} quantity(o) \cdot x_{ovc} \leq vehCapa \quad v \in V, \quad (6)$$

$$\sum_{v \in V} \sum_{c \in C} x_{ovc} = 1 \quad o \in O, \quad (7)$$

$$\sum_{o \in ordCust(j)} \sum_{c \in C} x_{ovc} \leq |O| \cdot \sum_{i \in L} b_{ijv} \quad v \in V, j \in L_{cust}, \quad (8)$$

$$\sum_{o \in O} x_{ovc} \leq |O| \cdot a_{vc} \quad c \in C, v \in V, \quad (9)$$

$$\sum_{o \in O_p} x_{ovc} \leq |O| \cdot (1 - x_{rv}) \quad v \in V, c \in C, \quad (10)$$

$$a_{vc} \in \{0, 1\} \quad v \in V, c \in C, \quad (11)$$

$$b_{ijv} \in \{0, 1\} \quad i, j \in L, v \in V, \quad (12)$$

$$u_{iv} \in \{1, \dots, |L|\} \quad i \in L, v \in V, \quad (13)$$

$$x_{ovc} \in \{0, 1\} \quad o \in O, v \in V, c \in C. \quad (14)$$

3.4.2 Decision variables

The model uses four groups of decision variables, the ranges of which are defined in equations (11) to (14).

The first is a binary variable relating to the route, which indicates whether a given vehicle travels directly from one to another. The second is an integer position variable, which indicates the position of a customer along a vehicle's route and serves to eliminate partial tours. The third is a binary assignment variable indicating whether a given order is delivered by a given vehicle to a given compartment. The fourth is a binary auxiliary variable indicating whether a compartment is active for a vehicle, where the number of active compartments for a vehicle is defined as the sum of this auxiliary variable across all the vehicle's compartments. In the remainder of this thesis, this quantity is defined as the number of compartments in the vehicle.

The assignment of orders to compartments already implies their activation, as indicated in Equation (9), while decisions regarding the order and allocation of orders are the same for any vehicle routing problem. What distinguishes the MCVRP_LU method is the allocation of orders to compartments, combined with the choice of the number of compartments (Hübner & Ostermeier, 2019).

3.4.3 Objective function

The first term of Equation (1) in brackets sums the loading costs of all vehicles, which depend on the number of compartments and, consequently, on the number of product segments transported per journey. The second term corresponds to the total transport costs, which consist of the weighted distance costs according to the selected product segments of the route. The third term corresponds to the total unloading cost, which is equal to the unloading fee per stop multiplied by the number of customer stops along the route, although the inner sum relates only to arrivals at customer locations, as the distribution centre (DC) is not counted as a stop. Given that the loading and unloading fees are adjusted according to the number of compartments, which in turn is determined by the decision-making variables, the objective is non-linear (Hübner & Ostermeier, 2019).

3.4.4 Constraints

The objective is minimised subject to the constraints given in equations (2) to (14).

Equations (2) and (3) ensure that each vehicle departs from the distribution centre at most once and that a vehicle arriving at a location will depart from it again, thereby ensuring the continuity of the flow of goods. Equations (4) and (5) formulate the Miller-Tucker-Zemlin model and

eliminate the partial tours by ordering the locations along a route and defining the distribution centre as the starting point. Equation (6) ensures that vehicle capacities are adhered to in all compartments. Equation (7) stipulates that each order must be assigned exactly once. Equation (8) links the visit to the load and requires a vehicle to visit a customer when an order from that customer is loaded onto the vehicle. Equation (9) activates a compartment each time an order is assigned to it. Equation (10) represents the incompatibility constraint and prevents orders of different product segments from sharing the same loading space. Equations (11) to (14) define the value ranges of the variables (Hübner & Ostermeier, 2019).

3.5 Solution methodology

The approach to the solution follows the two-stage framework proposed by Hübner and Ostermeier (2019), which is based on the study of operators by Derigs et al. (2011). In the first stage, an initial feasible solution is generated using the parallel savings algorithm of Clarke and Wright (1964), which provides fast and feasible initial solutions for VRP and MCVRP (Frank et al., 2021). In the second phase, this solution is improved via a large neighbourhood search (LNS), during which a set of orders is repeatedly pruned using a removal operator according to Shaw (Shaw, 1997) and reintroduced via the regret operator (Ropke & Pisinger, 2006). The acceptance of candidate solutions is governed by a ‘Record-to-Record Travel’ criterion (Dueck, 1993), whilst a control mechanism varies the search through periodic large-scale pruning and terminates it after a limited number of unsuccessful iterations. The remaining subsections describe these components in the order in which they are assembled.

3.5.1 Construction heuristic: parallel savings

The initial solution is generated based on the parallel version of the Clarke and Wright (1964) saving algorithm, which is a standard and fast heuristic construction method for multi-compartment routing problems (Hübner & Ostermeier, 2019; Muyldermans & Pang, 2010, cited in Hübner & Ostermeier, 2019). The parallel variant is chosen because it tends to provide better solutions for capacity-constrained routing compared with the sequential variant (Laporte, 2009, cited in Hübner & Ostermeier, 2019).

The process begins with a basic route for each order and combines the routes in a ‘greedy’ manner, in descending order of savings. The savings associated with combining two routes are

calculated as the distance from the final location of the first route to the distribution centre (DC), plus the distance from the DC to the first location of the second route, minus the direct distance between these two locations. Merging takes place only if the resulting route remains viable in terms of the vehicle's capacity as per equation (6) and the compartment's upper limit. To serve customers with multiple orders, orders from the same customer are grouped at a common stop prior to evaluating the savings, and any potential merge that involves visiting the same customer twice is rejected. The procedure is summarised in Figure 3.

Figure 3: Pseudocode for parallel savings construction (adapted from Hübner & Ostermeier, 2019)

Input: orders O , distance matrix, vehicle capacity, compartment ceiling

1. Create one route per customer, holding all their orders in a single stop (split off a new route if a stop would exceed $vehCapa$)
2. For each ordered pair (i, j) on different customers, compute $saving(i, j) = dist(cust_i, 0) + dist(0, cust_j) - dist(cust_i, cust_j)$
3. Sort candidate merges by decreasing saving
4. For each merge in order, if i is a route tail and j a route head, and the result respects $vehCapa$, $\bar{c}$, and one stop per customer, join the routes
5. Drop empty routes and return the initial solution

The construction provides a feasible starting point whose quality is then improved by the LNS. Its role is deliberately limited, since the parallel savings heuristic supplies fast feasible solutions and the burden of solution quality rests with the improvement stage (Hübner & Ostermeier, 2019).

3.5.2 Shaw removal

The removal algorithm follows Shaw's (1997) approach, which was chosen by Hübner and Ostermeier (2019), as it deletes orders with a similar structure, a key factor in this context, given that the combination of disparate segments increases the total cost and because it is one of the removal algorithms for the solution set proposed by Derigs et al. (2011). The operator removes a predetermined number of orders from the current solution, selecting the first order at random and each subsequent order based on a similarity measure and a controlled degree of randomness.

The set of all orders is divided into removed orders and matched orders; the two subsets are disjoint and, together, cover the entire set of orders. Once the first order has been randomly

deleted, the similarity between an order that has already been deleted and a candidate order that remains matched is calculated using the relatedness measure in equation (15).

$$\frac{1}{R_{os}} := \phi \cdot \frac{cost_{os}}{cost_{max}} + \omega \cdot prod_{os} + \psi \cdot \frac{|quantity(o) - quantity(s)|}{quantity_{max}}. \quad (15)$$

A lower relatedness index value indicates greater similarity. The first term in equation (15) represents the transport costs between the customers placing the two orders, normalised with respect to the maximum transport costs between any two customers and reflects both proximity in terms of transport and unloading, as orders from the same or neighbouring customers are advantageous. The second term is the product segment index defined in equation (16), which takes the value one when the two orders belong to different segments and the value zero in all other cases; thus, two orders from the same segment that can share storage space are considered more similar and the loading effort associated with mixing segments is implicitly taken into account. The third term is the normalised difference in order quantity, as exchanging orders of similar size helps to ensure feasibility. The three weighting coefficients express the relative importance of distance, segment and size (Hübner and Ostermeier, 2019).

$$prod_{os} := \begin{cases} 1 & \text{if } segment(o) \neq segment(s), \\ 0 & \text{otherwise.} \end{cases} \quad (16)$$

This relatedness model has now become a fundamental element of research in the field of multi-compartment models. Frank et al. (2021) use the same three-part measure (transport cost, segment identity and shipment size) in the route planning phase of their multi-temporal framework, while Chen et al. (2019) use a modified version of Shaw's algorithm for the same family of problems, adapting the relatedness measure to take time windows into account. The operator is a standard, well-documented component, not unique to the MCVRP_LU model. This benefits reconstruction, as its behaviour can be verified against broader literature instead of relying solely on a single source.

The selection is carried out by sorting the candidate orders in descending order of similarity, so that the most similar orders are ranked first, and then selecting the orders that lie a random number of positions further down the ranking. This position is obtained by drawing a uniformly

distributed random number from the unit interval and raising it to the power of the ranking exponent, which determines the balance between intensification and diversification. A value of one for the exponent limits the selection to a uniformly distributed random sample, whilst higher values steer the selection towards the most similar orders; the resulting position is rounded to the nearest integer if it is not an integer (Hübner & Ostermeier, 2019). The operator is summarised in Figure 4.

Figure 4: Pseudocode for Shaw removal (adapted from Hübner & Ostermeier, 2019)

Shaw removal **Input:** number of orders to remove, the three relatedness weights, the ranking exponent

1. Set $\text{removed} \leftarrow \emptyset$ and $\text{assigned} \leftarrow$ all orders in the current solution
2. Select one order uniformly at random from assigned, move it to removed
3. While $|\text{removed}| < q$ do
4. Select a reference order ref uniformly at random from removed
5. For each order $o \in \text{assigned}$ do
6. compute $1/R(\text{ref}, o) = \phi \cdot \text{cost}(\text{ref}, o)/\text{cost_max} + \omega \cdot \text{prod}(\text{ref}, o) + \psi \cdot |\text{quantity}(o) - \text{quantity}(\text{ref})|/\text{quantity_max}$,
where $\text{prod}(\text{ref}, o) = 1$ if the two orders belong to different product segments and 0 otherwise
7. End for
8. Sort assigned by ascending $1/R(\text{ref}, o)$, so that the most similar order takes rank 0
9. Draw z uniformly from $[0, 1)$ and set $\text{position} \leftarrow \lfloor z^\alpha \cdot |\text{assigned}| \rfloor$
10. Move the order at that position from assigned to removed
11. End while
12. Remove the orders in removed from their routes and delete any route left without orders
13. Return the reduced solution and removed

3.5.3 Regret-k insertion

The reintegration operator follows the regret-based approach of Ropke and Pisinger (2006), which has been adapted so that the integration costs reflect the overall objective of the three terms in Equation (1), rather than just the transfer costs. For each removed order, the regret value of a given degree measures how much is lost if the order is not placed at its current optimal location, whereby the differences between the lowest-cost location and the next-lowest-cost locations along the various routes are summed. For a degree of regret of at least two, the regret value is given by Equation (17).

$$\text{regret}_k := \sum_{u=2}^k (\text{totalcost}_{v_u} - \text{totalcost}_{v_1}). \quad (17)$$

In Equation (17), the routes are ranked according to the total cost involved in assigning the order to each route, so that the route at the top of the ranking is the most economical. In each iteration, the order with the highest regret value is assigned to the route best suited to it;

thereafter, the regret values of the remaining orders that had been eliminated are recalculated, as the available insertion points may have changed. The process is repeated until all the orders that had been eliminated have been reintroduced (Hübner & Ostermeier, 2019).

The cost of a single insertion is calculated as the change in the total objective in Equation (1), that is, as the combined change in the loading, transport and unloading costs. An order can be incorporated in two ways. If the relevant customer has already been visited during the proposed route, the order can be added to the existing stop, which entails no additional stop at the customer's premises and, therefore, incurs no additional general unloading costs; however, it may trigger a new load compartment if the order introduces a new product segment to the vehicle's route, which increases loading costs and, potentially, the unloading fee. Alternatively, the order may result in a new stop, which adds the corresponding detour and general unloading costs. The most economically feasible option is selected and only entries that satisfy the vehicle's capacity as per Equation (6), the maximum capacity of the loading space and the incompatibility condition between product segments as per Equation (10) are accepted. The operator is summarised in Figure 5.

Figure 5: Pseudocode for Regret-k insertion (adapted from Hübner & Ostermeier, 2019)

Input: removed order set, regret degree k (here $k = 2$)

1. **While orders remain to insert:**
2. **For each removed order o , evaluate with the full objective (1): the cheapest feasible insertion into each route (existing stop or new stop, subject to (6), $\bar{c}$, and (10)), plus opening a new route**
3. **Sort the options by ascending cost and compute the regret value (17) as the sum over the $k-1$ next-best options minus the cheapest**
4. **Insert the order with the highest regret at its cheapest option (existing stop, new stop, or new route)**
5. **Return the repaired solution**

By assessing the location not only on the basis of immediate costs, but also future costs, the regret operator defers orders for which good location options are most likely to disappear and it is precisely this characteristic that makes it so effective for this particular problem (Ropke & Pisinger, 2006; Hübner & Ostermeier, 2019). Frank et al. (2021) reach the same conclusion, starting from an opposite perspective: they describe the regret insertion as essential to the quality of the solution in this class of problems and identify it as the single factor requiring the greatest computational effort to implement.

3.5.4 Record-to-Record Travel acceptance

To determine whether a candidate solution, through removal and reinsertion, replaces the current solution, the ‘Record-to-Record Travel’ criterion (Dueck, 1993) is used. This is the acceptance rule applied by Hübner and Ostermeier (2019) in their LNS method. At any given time, two solutions are maintained: the current working solution, which serves as the starting point for the next cycle of destruction and repair, and the best solution found so far, referred to as the ‘Record’. Each new candidate solution is compared with the ‘Record’ rather than with the current working solution.

The rule allows for a certain degree of deterioration in exchange for the possibility of deviating from the optimal points. A candidate solution is accepted as a new provisional solution whenever its total cost does not exceed the previous record value by more than $D\%$. The value of D is small, so that the search remains close to the range of optimal costs achieved; it is not equal to zero, so the search is not forced to accept only strict improvements, which would lead it into a dead end as soon as there were no further possibilities for improvement in the immediate vicinity. If a candidate solution improves on the record, it simultaneously becomes the new record and the valid solution.

The deviation parameter D has been set at 0.9% , the value established by Hübner and Ostermeier (2019). Two characteristics of this acceptance rule are particularly worth highlighting. It is deterministic, as a candidate with a given cost is either accepted or rejected without any additional random selection, which means that a specific execution cycle is exactly reproducible from its random seed value. Furthermore, it depends on the record and not on the provisional solution; thus, as soon as a good solution is found, the acceptance range is defined around it and the search shifts towards increasingly narrow ranges as the record improves.

3.5.5 Parameter setting

The algorithm’s parameters correspond to the values specified by Hübner and Ostermeier (2019), which are summarised in Table 1. They are adopted unchanged in all runs, without any instance-specific adjustments.

Table 1: Parameters of the LNS (Hübner & Ostermeier, 2019)

Parameter	Symbol	Value	Purpose
Shaw distance weight	ϕ	0.4	weight of spatial proximity
Shaw segment weight	ω	0.4	weight of shared product segment
Shaw quantity weight	ψ	0.2	weight of similar order size
Ranking exponent	α	4	strength of the bias toward the most related order in Shaw removal
Regret degree	k	2	number of best insertion positions compared in regret insertion
RRT deviation	D	0.9%	acceptance band above the record
Reset border	-	500	idle iterations before the working solution is reset to the record
Termination limit	-	2000	idle iterations before the run ends

The adoption of the published set of parameters without independent fine-tuning follows a different logic to that usually applied when presenting a new algorithm. Eshtehadi, Demir and Huang (2020), for example, optimise the parameters of their adaptive neighbourhood search extension using a subset of sixteen of the eighty instances, running each candidate configuration three times and selecting the combination with the lowest mean deviation from the best value found across all configurations. This procedure is appropriate when the aim is to improve upon a published result. However, in a reproducibility study, an independently optimised set of parameters would introduce a second source of deviation, in addition to the reconstruction itself, as any deviation from the published results could be attributed to either the new implementation or the choice of parameters. Retaining the published values, as has been done in this case, isolates the reconstruction as the sole variable to be verified.

4. Results

In this chapter, the reconstructed solution algorithm is applied to the instances published by Hübner and Ostermeier (2019) and the results are compared with the reference solutions provided by the authors for the same instances. Section 4.1 describes the set of instances, the cost levels and the comparison protocol. Section 4.2 presents the total costs compared with the published reference values. Section 4.3 compares the structure of the solutions, namely the number of routes, the number of stops and the average number of active compartments per route. Section 4.4 compares the cost composition for the sole configuration for which the paper

publishes a breakdown. Section 4.5 illustrates the execution times and explains why the two implementations do not coincide in terms of computational effort. Section 4.6 examines the stability of the reconstruction across independent runs.

4.1 Setup and comparison

The reconstructed LNS algorithm was implemented in C++ and executed on a desktop system equipped with an AMD Ryzen 7 5700X3D processor (3.0 GHz) and 16 GB of DDR4 RAM, running on 64-bit Windows. The comparison set comprises the seven instances for which the authors publish both the instance data and the corresponding solution files. Table 2 summarises their characteristics.

Table 2: Test instances and setup (Hübner & Ostermeier, 2019; repository available at <https://www1.ku.de/wwf/pw/forschung/>)

Instance (file)	Short label	Test	Customers	Orders	Orders per cust.	Ø order size (TU)	Distance (to Customer)	Runs per cost level
199n_400o_9tu_3s	A.2, 9 TU	A.2	199	400	1 to 3	9	11 to 80	5 seeds
199n_400o_7tu_3s	A.2, 7 TU	A.1/A.2	199	400	1 to 3	7	11 to 80	5 seeds
199n_400o_5tu_3s	A.2, 5 TU	A.2	199	400	1 to 3	5	11 to 80	5 seeds
199n_400o_3tu_3s	A.2, 3 TU	A.2	199	400	1 to 3	3	11 to 80	5 seeds
199n_400o_7tu_3s	B	B	199	400	1 to 3	7	9 to 40	5 seeds
106n_198o_7tu_3s	C.1	C.1	106	198	1 to 3	7	9 to 40	5 seeds
106n_106o_7tu_3s	C.2	C.2	106	106	1	7	9 to 40	5 seeds

All scenarios are based on three product segments and a transport capacity of 33 TU. The figure given in the scenario name indicates the average order size, although quantities range from one TU to approximately twice the average, in accordance with the generation procedure described by Hübner and Ostermeier (2019). Each scenario is calculated at three handling cost levels, labelled ‘standard’, ‘+50%’ and ‘+100%’, where the increase relates exclusively to loading and unloading costs, whilst transport costs remain unchanged. For each scenario and each level, five independent runs are carried out using different random seeds, resulting in a total of 105 runs.

The reference values used in this chapter are taken from two sources. The first source is the solution value indicated in each published solution file, which is used for the comparison of total costs in section 4.2. The second source is the information on the different routes stored in those same files, that is, the number of routes, the number of stops and the number of active compartments, which is used for the structural comparison in section 4.3. Whenever a piece of data is included directly in the article itself, such as the cost percentages in Table 14 and the structure of the solution for Test B in Table 15, that data is reproduced exactly as given. No reference values in this chapter are derived anew from the published solutions and each published piece of data is taken into account exactly as it appears in the solution file or in the article.

The values given for the reconstruction are the total costs and solution structures generated by the solver itself using the cost parameters documented in section 3.3. The best value for a group refers to the lowest total cost amongst the five runs for this instance and at this cost level and the average value refers to the arithmetic mean of those five runs.

4.2 Total cost against the published reference value

Tables 3 to 5 present, for each cost level, the best and average total costs derived from the five runs of the reconstruction, together with the reference value and the resulting deviation. The deviation is calculated as $(\text{own best} - \text{published}) / \text{published}$.

Table 3: Total costs at standard handling costs

Instance	Own best	Own Ø	Published	Deviation
A.2, 9 TU	13949.68	14044.02	15127.82	-7.8%
A.2, 7 TU	11105.89	11162.73	12277.84	-9.5%
A.2, 5 TU	8197.43	8227.33	9029.32	-9.2%
A.2, 3 TU	6006.18	6018.71	6613.96	-9.2%
B	9188.27	9200.94	9938.63	-7.6%
C.1	4820.25	4835.57	5134.75	-6.1%
C.2	2877.22	2912.09	3025.00	-4.9%

Table 4: Total cost at +50% handling cost

Instance	Own best	Own Ø	Published	Deviation
A.2, 9 TU	14686.43	14750.20	15782.16	−6.9%
A.2, 7 TU	11671.74	11756.08	12785.00	−8.7%
A.2, 5 TU	8695.20	8718.66	9549.04	−8.9%
A.2, 3 TU	6387.77	6397.77	7054.97	−9.5%
B	9756.41	9784.20	10599.18	−8.0%
C.1	5126.01	5147.97	5445.21	−5.9%
C.2	3042.29	3072.27	3068.60	−0.9%

Table 5: Total cost at +100% handling cost

Instance	Own best	Own Ø	Published	Deviation
A.2, 9 TU	15336.74	15402.58	16220.68	−5.5%
A.2, 7 TU	12228.77	12346.86	13346.46	−8.4%
A.2, 5 TU	9175.32	9191.88	10108.45	−9.2%
A.2, 3 TU	6771.36	6781.56	7507.38	−9.8%
B	10357.44	10375.69	11042.39	−6.2%
C.1	5414.90	5444.64	5706.30	−5.1%
C.2	3183.04	3237.26	3063.18	+3.9%

Out of a total of 21 comparisons, the reconstruction shows, in 20 cases, total costs lower than the reference value, with deviations ranging mostly between −4.9% and −9.8%, while in only one case, specifically case C.2, are the total costs higher, specifically by +3.9% compared with the reference value of +100%. The deviations are greater in the cases with 199 customers and 400 orders, ranging between −5.4% and −9.8%, whilst in the two scenarios with 106 customers they are smaller, ranging between −0.9% and −6.1%, excluding the single positive value.

The difference between the best and mean columns within the same reconstruction is minimal and, in each group, is below 1% in 19 of 21 groups, at most 1.7%, which suggests that the best value for a group is representative and does not constitute an outlier in a highly scattered sample.

This thesis reproduces the comparison as observed and does not aim to determine the cause of the difference between the two data sets. The published values are reproduced as they appear in the solution files, whilst reconstructing the precise evaluation on which they are based would require assumptions that go beyond the scope of this thesis and the attached files, primarily because the underlying loading and unloading measurements are derived from a time-and-motion study carried out with the retailer that is the subject of the case study and are subject to a confidentiality agreement. Therefore, Section 4.3 focuses on the structure of the solution, which is explicitly described in both sources and allows for a direct comparison.

4.3 Solution structure

Tables 6 to 8 compare the structure of the reconstructed best solution with that of the published solution for each scenario and each cost level. The number of routes, the number of stops and the average number of active compartments per route are derived directly from the route information contained in both sources.

Table 6: Solution structure at standard handling costs

Instance	Own routes	Own stops	Own Ø comp.	Pub. routes	Pub. stops	Pub. Ø comp.
A.2, 9 TU	124	242	2.34	120	285	2.19
A.2, 7 TU	96	223	2.72	93	258	2.44
A.2, 5 TU	64	211	2.97	64	232	2.84
A.2, 3 TU	39	199	3.00	40	213	2.92
B	91	228	2.65	91	256	2.47
C.1	47	120	2.64	47	124	2.53
C.2	23	106	2.35	24	106	1.71

Table 7: Solution structure at +50% handling costs

Instance	Own routes	Own stops	Own Ø comp.	Pub. routes	Pub. stops	Pub. Ø comp.
A.2, 9 TU	123	246	2.34	120	318	1.83
A.2, 7 TU	95	224	2.69	93	263	2.37
A.2, 5 TU	64	203	2.95	64	228	2.81
A.2, 3 TU	39	199	3.00	40	206	2.95
B	92	230	2.55	91	276	2.17
C.1	46	127	2.57	48	125	2.44
C.2	23	106	2.22	24	106	1.25

Table 8: Solution structure at +100% handling costs

Instance	Own routes	Own stops	Own Ø comp.	Pub. routes	Pub. stops	Pub. Ø comp.
A.2, 9 TU	122	244	2.39	120	323	1.73
A.2, 7 TU	93	232	2.65	93	280	2.14
A.2, 5 TU	64	204	2.94	65	234	2.69
A.2, 3 TU	39	199	3.00	40	206	2.98
B	95	224	2.64	91	327	1.55
C.1	47	119	2.64	47	143	1.98
C.2	23	106	1.61	24	106	1.17

With standard handling costs, the two solution structures are similar. The number of routes agrees to within four routes in all seven cases and is identical in three of them, while the average number of active compartments differs by between 0.08 and 0.28, with the exception of case C.2, where the reconstruction uses 2.35 compartments per route, compared with 1.71 in the published solution. In six of the seven cases, the reconstruction requires fewer stops, between 4 and 43 fewer, which corresponds to a solution in which several segments for the same customer are delivered simultaneously. In C.2, the two solutions agree at 106 stops, as each customer places exactly one order and 106 therefore represents a lower bound.

The two structures react differently to an increase in handling costs. In the published solutions, the average number of active compartments decreases significantly, which is particularly evident in Test B, where it changes from the typical value of 2.47 to 2.17 with a 50% increase and to 1.55 with a 100% increase, while at the same time, the number of stops increases from 256 to 276 and then to 327. The same trend is observed in the case with an average order size of nine transport units, where the published value decreases from 2.19 to 1.73, as well as in case C.1, where it decreases from 2.53 to 1.98. In the reconstruction, the corresponding deviation is significantly smaller. In Test B, the average number of compartments remains between 2.55 and 2.65 across all three stages, whilst the number of stops ranges between 224 and 230.

Consequently, the reconstruction faithfully reproduces the structure of the published solutions with standard handling costs, while the trend towards vehicles with fewer cargo compartments, which is observed in the published solutions with higher handling costs, is only evident to a limited extent in the reconstruction. The most striking exception is case C.2, at the +100% level, where the reconstruction also significantly reduces the average number of compartments, from 2.35 at the standard level to 1.61.

4.4 Cost composition

In Table 14, Hübner and Ostermeier (2019) give the proportion of the costs of loading, transport and unloading relative to the total costs for Test B. In Table 9, these values are compared with the corresponding proportions in the reconstructed best solution for the same instance and the same cost level.

Table 9: Cost shares for Test B at standard handling costs

Source	Loading	Transportation	Unloading
Own reconstruction	7.0%	87.3%	5.7%
Hübner and Ostermeier (2019), Table 14	6.1%	87.6%	6.3%

The two datasets show a close correlation, with the maximum difference in each category amounting to 0.9 percentage points. In both cases, transport is the most significant category, accounting for approximately 87% of total costs, whilst loading and unloading operations together account for between 12% and 13%, which is consistent with the findings in section 2 of the article, according to which these operations can account for up to 20% of total costs. The cost structure of the model applied here therefore corresponds to the cost structure reported for the original application.

This comparison is limited to Test B at the level of standardised costs, as this is the configuration for which the article publishes a cost analysis. Although the published solution files contain the route structure, they do not include a cost breakdown by component; consequently, it is not possible to supplement Table 9 with further examples based on the available sources.

4.5 Runtime

Table 10 reports the average runtime per run of the reconstruction, together with the runtime figures published in the article for comparable configurations.

Table 10: Average runtime per run in seconds

Instance	Own, Std	Own, +50%	Own, +100%	Paper reference
A.2, 9 TU	1566	1526	1551	24
A.2, 7 TU	1771	1745	1725	41
A.2, 5 TU	2252	2134	2229	n.a.
A.2, 3 TU	3090	3024	3049	168
B	1514	1591	1496	n.a.
C.1	236	234	216	25 to 26
C.2	66	71	76	n.a.

The qualitative pattern reported in the article is reproduced in full. The average order size is the decisive factor influencing runtime, which increases monotonically as order size decreases: from 1566 seconds with nine transport units to 3090 seconds with three transport units in cases involving 199 customers. This series corresponds to that mentioned in the article and the explanation provided there is also valid in this case, namely that smaller orders allow a greater number of orders to be loaded onto a single vehicle, which means there are more insertion positions available at each repair stage. The size of the scenarios also points in the same direction: the two scenarios with 106 customers are an order of magnitude faster than those with 199 customers. The variation in runtime across the three handling cost levels stays below 7% for the five instances with 199 customers and up to 15% for the smallest.

The actual execution times of the reconstruction far exceed the published figures, by a factor of 9 to 65 for similar configurations. As the implementation is in C++ and runs on significantly more modern hardware than the 1.8 GHz computer used in the original study, such a large difference cannot be explained by processor speed alone. The decisive factor is likely to be the number of search operations performed per execution. Hübner and Ostermeier (2019) set their termination condition at 2000 unsuccessful iterations with a restart limit of 500. They mention that the number of orders removed is sampled uniformly within instance-dependent bounds, but these bounds are specified only for the runtime instances from Section 6.2.2 of the article, not for the instances of Tests A to C that are of interest to us here. It also remains unclear whether the regret evaluation is restricted to a subset of the candidate tours. Both decisions influence the number of insertion evaluations per iteration and thus the total cost of an execution that terminates after the same number of unsuccessful iterations.

The execution times listed in Table 10 show that the reconstruction runs and the published runs do not correspond in terms of search effort, meaning that the total costs compared in Section 4.2 are not based on equivalent computational effort. A search that performs an order of magnitude or more as many insertion evaluations before aborting would likely achieve lower target values for the same instances and the difference stated in Section 4.2 falls within the range that such an advantage could produce. Therefore, the observed difference is presented as it stands, without claiming that it is attributable to the target evaluation rather than the scope of the search. To disentangle these factors, either the scaling rule of the original implementation, which is not specified in the documentation, would need to be obtained or comparable runs would need to be carried out using the published runtimes; however, this lies outside the scope of the protocol described in Section 4.1.

This limitation does not apply exclusively to the present comparison. Coelho, Cordeau and Laporte (2012) point out that, in a heuristic applied to a problem relating to route calculation for multiple vehicles, approximately 65% of the total execution time corresponds to a sub-problem that is solved repeatedly and a further 7% to a solution optimisation phase; a comparison of execution times between different implementations is only meaningful if this breakdown is available for both parts, however, it is precisely this information that is missing in the present case.

4.6 Stability across runs

Table 11 reports the coefficient of variation of total costs across the five runs performed for each instance and cost level.

Table 11: Coefficient of variation of total costs across five runs

Instance	Std	+50%	+100%
A.2, 9 TU	0.42%	0.27%	0.32%
A.2, 7 TU	0.47%	0.46%	0.57%
A.2, 5 TU	0.31%	0.21%	0.18%
A.2, 3 TU	0.13%	0.15%	0.14%
B	0.14%	0.21%	0.13%
C.1	0.29%	0.44%	0.62%
C.2	1.12%	0.81%	1.25%

Overall, the dispersion is modest and, across the 21 groups, ranges between 0.13% and 1.25%, while the difference between the best and worst results within a group never exceeds 3.0%. Consequently, the search implemented in the reconstruction is stable, and the best value across the five runs is close to the average of these five runs, as can already be seen in the first two columns of Tables 3 to 5.

This measure of dispersion is consistent with the values reported in other studies for comparable meta-heuristics; Coelho, Cordeau and Laporte (2012) apply the same coefficient of variation to a metaheuristic for a related routing problem and report an average value of approximately 0.01 (i.e. around 1%) in repeated runs, which is of the same order of magnitude as the values obtained here.

The dispersion is greatest in instance C.2, the smallest, with 106 orders and exactly one order per customer, and is smallest in instances with small average order sizes, where the capacity constraint leaves relatively little structural freedom and almost all vehicles are used with the maximum of three compartments. The use of five runs per configuration rather than the 100 runs mentioned in the article was a deliberate scope decision and the low dispersion observed indicates that the resulting loss of precision in the best values presented is minimal compared with the differences analysed in this chapter.

5. Conclusion

5.1 Summary of the study

This thesis treated the multi-compartment vehicle routing problem with loading and unloading costs, as described by Hübner and Ostermeier (2019), as a specification to be implemented rather than as a result to be cited. The model was reformulated to explicitly state its assumptions and the two-stage solution method, a parallel savings construction followed by a large neighbourhood search with Shaw removal, Regret-k insertion and Record-to-Record Travel acceptance, was implemented in C++ based solely on the published description and the set of parameters. The solver was applied to the seven instances for which the authors provide both the instance data and the solution files, with three levels of handling costs and five independent runs for each, resulting in a total of 105 runs. The comparison was carried out across four dimensions: total cost, solution structure, cost composition and execution time, so that reproducibility would not be assessed solely on the basis of a single overall figure.

5.2 Answers to the research questions

The first research question concerned whether a solver created from the published description would reproduce the published benchmark results, both numerically and structurally. From a structural point of view, the reconstruction faithfully reproduces the published solutions with standard operating costs, such that the number of routes across the seven instances matches within a tolerance of four routes and the average number of compartments, except in instance C.2, varies between 0.08 and 0.28, while at the same time, a smaller number of stops is consistently generated. The trend towards vehicles with fewer compartments, as shown by the published solutions with higher handling costs, is only evident to a limited extent in the reconstruction. As regards the breakdown of costs, the percentages calculated for Test B relating to loading, transport and unloading correspond to the values given in the article, with the maximum deviation for each component being 0.9 percentage points. In numerical terms, the total cost of the reconstruction lies below the published reference values in 20 of the 21 comparisons, with deviations ranging mostly from -4.9% to -9.8% , and only in one comparison does it exceed these values. This difference is reported as observed. Section 4.5 shows that the two implementations are not matched in search effort, so it cannot be assigned to a single cause on the evidence available here.

The second research question concerned those aspects of the published description that leave room for interpretation in the context of an independent implementation. Two areas emerged from the experiments. The first relates to cost parameters, which are given in relative terms in the article, as the underlying measurements are confidential; consequently, an implementation must determine its exact application independently. The second area concerns the algorithmic control parameters, where, although the termination threshold and the restart threshold are established, neither the number of failed iterations nor the sampling of the number of deleted orders is specified; consequently, the search effort for each run, and, therefore, also the running time, remains undetermined. Neither of these two observations constitutes a shortcoming of the article, whose cost data are based on confidential empirical measurements and whose algorithmic description is inevitably condensed within the space available in a journal publication. However, both are issues on which a replication study must make a decision and, therefore, it is worth mentioning them explicitly in studies of this kind.

Appendix

List of Figures

		Page
Figure 1	Load structure and delivery route of a multi-compartment vehicle	8
Figure 2	Assignment of segment specific orders to two vehicle tours	11
Figure 3	Pseudocode for parallel savings construction	18
Figure 4	Pseudocode for Shaw removal	20
Figure 5	Pseudocode for regret-k insertion	21

List of Tables

		Page
Table 1	Parameters of the LNS	23
Table 2	Test instances and setup	24
Table 3	Total cost at standard handling costs	26
Table 4	Total cost at +50% handling costs	26
Table 5	Total cost at +100% handling costs	27
Table 6	Solution structure at standard handling costs	28
Table 7	Solution structure at +50% handling costs	28
Table 8	Solution structure at +100% handling costs	29
Table 9	Cost shares for Test B at standard handling costs	30
Table 10	Average runtime per run in seconds	31
Table 11	Coefficient of variation of total costs across five runs	32

List of Abbreviations

CVRP Capacitated vehicle routing problem

DC Distribution centre

LNS Large neighbourhood search

MCV Multi-compartment vehicle

MCVRP Multi-compartment vehicle routing problem

MCVRP_LU Multi-compartment vehicle routing problem with loading and unloading costs

RRT Record-to-Record Travel

SCV Single-compartment vehicle

TU Transport unit

VRP Vehicle routing problem

Results for standard, 50% and 100%

199n_400o_9tu_3s (TestA.2) - Standard

Seed	Loading cost	Traveling cost	Unloading cost	Total cost	Stops	Routes	Average Compartment
1	768.87	12705	574.16	14048.03	252	121	2.388
2	809.12	12742	521.27	14072.39	228	125	2.432
3	832.83	12767	507.74	14107.57	222	126	2.484
4	722.25	12675	552.43	13949.68	242	124	2.339
5	787.68	12712	542.75	14042.43	238	123	2.366

199n_400o_9tu_3s (TestA.2) +50%

Seed	Loading cost	Traveling cost	Unloading cost	Total cost	Stops	Routes	Average Compartment
1	1249.25	12767	761.61	14777.86	222	126	2.484
2	1189.81	12756	791.5	14737.31	231	125	2.384
3	1249.25	12767	761.61	14777.86	222	126	2.484
4	1237.31	12763	771.21	14771.52	225	126	2.460
5	1150.11	12695	841.32	14686.43	246	123	2.341

199n_400o_9tu_3s (TestA.2) +100%

Seed	Loading cost	Traveling cost	Unloading cost	Total cost	Stops	Routes	Average Compartment
1	1665.66	12767	1015.49	15448.15	222	126	2.508
2	1443.21	12784	1143.9	15371.11	252	119	2.277
3	1665.66	12767	1015.49	15448.15	222	126	2.484
4	1465.26	12732	1211.47	15408.73	265	119	2.311
5	1549.7	12672	1115.04	15336.74	244	122	2.385

199n_400o_5tu_3s (TestA.2) Standard

Seed	Loading cost	Traveling cost	Unloading cost	Total cost	Stops	Routes	Average Compartment
1	501.38	7294	461.49	8256.87	201	64	2.953
2	501.38	7234	470.68	8206.06	205	64	2.953
3	501.38	7280	463.79	8245.17	202	64	2.953
4	503.98	7209	484.45	8197.43	211	64	2.969
5	493.58	7260	477.56	8231.14	208	64	2.906

199n_400o_5tu_3s (TestA.2) +50%

Seed	Loading cost	Traveling cost	Unloading cost	Total cost	Stops	Routes	Average Compartment
1	752.07	7297	688.8	8737.87	200	64	2.953
2	752.07	7244	699.13	8695.2	203	64	2.953
3	752.07	7283	692.24	8727.31	201	64	2.953
4	752.07	7246	706.02	8704.09	205	64	2.953
5	744.27	7282	702.57	8728.84	204	64	2.922

199n_400o_5tu_3s (TestA.2) +100%

Seed	Loading cost	Traveling cost	Unloading cost	Total cost	Stops	Routes	Average Compartment
1	997.56	7262	955.13	9214.69	208	64	2.938
2	997.56	7241	936.76	9175.32	204	64	2.938
3	997.56	7260	936.76	9194.32	204	64	2.938
4	1002.77	7263	932.17	9197.94	203	64	2.953
5	1002.77	7233	941.36	9177.13	205	64	2.953

199n_400o_7tu_3s (TestB) Standard

Seed	Loading cost	Traveling cost	Unloading cost	Total cost	Stops	Routes	Average Compartment
1	643.39	8046	510.53	9199.92	223	93	2.602
2	640.79	8024	523.48	9188.27	228	91	2.648
3	669.7	8043	490.11	9202.81	214	95	2.653
4	688.06	8057	477.07	9222.13	208	95	2.726
5	595.81	8024	571.78	9191.59	250	89	2.517

199n_400o_7tu_3s (TestB) 50%

Seed	Loading cost	Traveling cost	Unloading cost	Total cost	Stops	Routes	Average Compartment
1	961.19	8042	772.68	9775.87	225	93	2.591
2	1032.1	8057	715.61	9804.71	208	95	2.726
3	980.92	8046	752.39	9779.31	219	94	2.617
4	1032.1	8057	715.61	9804.71	208	95	2.726
5	937.09	8032	787.32	9756.41	230	92	2.554

199n_400o_7tu_3s (TestB) 100%

Seed	Loading cost	Traveling cost	Unloading cost	Total cost	Stops	Routes	Average Compartment
1	1265.68	8060	1042.06	10367.74	228	94	2.532
2	1376.13	8057	954.15	10387.28	208	95	2.726
3	1334.22	8043	980.22	10357.44	224	95	2.642
4	1376.13	8057	954.15	10387.28	208	95	2.726
5	1265.37	8057	1056.32	10378.69	231	93	2.559

199n_400o_3tu_3s (TestA.2) Standard

Seed	Loading cost	Traveling cost	Unloading cost	Total cost	Stops	Routes	Average Compartment
1	310.28	5257	456.9	6024.18	199	39	3.000
2	310.28	5257	456.9	6024.18	199	39	3.000
3	310.28	5257	456.95	6024.23	199	39	3.000
4	312.88	5245	456.9	6014.78	199	39	3.000
5	310.28	5239	456.9	6006.18	199	39	3.000

199n_400o_3tu_3s (TestA.2) 50%

Seed	Loading cost	Traveling cost	Unloading cost	Total cost	Stops	Routes	Average Compartment
1	465.42	5257	685.35	6407.77	199	39	3.000
2	465.42	5257	685.35	6407.77	199	39	3.000
3	465.42	5245	685.35	6395.77	199	39	3.000
4	465.42	5237	685.35	6387.77	199	39	3.000
5	465.42	5239	685.35	6389.77	199	39	3.000

199n_400o_3tu_3s (TestA.2) 100%

Seed	Loading cost	Traveling cost	Unloading cost	Total cost	Stops	Routes	Average Compartment
1	620.56	5257	913.8	6791.36	199	39	3.000
2	620.56	5257	913.8	6791.36	199	39	3.000
3	620.56	5245	913.8	6779.36	199	39	3.000
4	620.56	5237	913.8	6771.36	199	39	3.000
5	620.56	5240	913.8	6774.36	199	39	3.000

199n_400o_7tu_3s (TestA.2) Standard

Seed	Loading cost	Traveling cost	Unloading cost	Total cost	Stops	Routes	Average Compartment
1	685.62	9909	511.27	11105.89	223	96	2.719
2	640.64	9904	561.53	11106.17	245	92	2.620
3	717.13	10020	477.56	11214.69	208	97	2.784
4	688.22	10005	499.79	11193.01	218	96	2.698
5	701.53	10001	491.34	11193.87	214	97	2.722

199n_400o_7tu_3s (TestA.2) 50%

Seed	Loading cost	Traveling cost	Unloading cost	Total cost	Stops	Routes	Average Compartment
1	1020.4	9881	770.34	11671.74	224	95	2.695
2	984.36	9975	794.94	11754.3	232	95	2.600
3	1075.7	10020	716.35	11812.05	208	97	2.784
4	1047.93	10003	742.42	11793.35	216	97	2.711
5	972.89	9982	794.08	11748.97	231	93	2.624

199n_400o_7tu_3s (TestA.2) 100%

Seed	Loading cost	Traveling cost	Unloading cost	Total cost	Stops	Routes	Average Compartment
1	1307.9	9857	1063.87	12228.77	232	93	2.645
2	1328.39	9997	1051.24	12376.63	230	94	2.660
3	1434.26	10020	955.13	12409.39	208	97	2.784
4	1376.44	9964	999.58	12340.02	218	96	2.698
5	1318.62	9997	1063.87	12379.49	232	94	2.638

106n_198o_7tu_3s (TestC.1) Standard

Seed	Loading cost	Traveling cost	Unloading cost	Total cost	Stops	Routes	Average Compartment
1	348.08	4221	273.22	4842.3	119	48	2.729
2	329.73	4215	275.52	4820.25	120	47	2.638
3	348.24	4218	261.74	4827.98	114	49	2.673
4	342.73	4243	270.19	4855.92	118	48	2.688
5	334.62	4222	274.78	4831.4	120	46	2.739

106n_198o_7tu_3s (TestC.1) 50%

Seed	Loading cost	Traveling cost	Unloading cost	Total cost	Stops	Routes	Average Compartment
1	474.63	4262	417.95	5154.58	122	47	2.532
2	541.86	4258	382.28	5182.14	111	49	2.776
3	470.73	4219	436.28	5126.01	127	46	2.565
4	478.76	4251	419.06	5148.82	122	47	2.553
5	474.63	4225	428.65	5128.28	125	47	2.532

106n_198o_7tu_3s (TestC.1) 100%

Seed	Loading cost	Traveling cost	Unloading cost	Total cost	Stops	Routes	Average Compartment
1	675.06	4275	538.9	5488.96	118	49	2.592
2	659.46	4209	546.44	5414.9	119	47	2.638
3	696.17	4207	514.3	5417.47	112	48	2.729
4	664.35	4221	544.97	5430.32	119	47	2.660
5	669.55	4259	543	5471.55	119	48	2.625

106n_106o_7tu_3s (TestC.2) Standard

Seed	Loading cost	Traveling cost	Unloading cost	Total cost	Stops	Routes	Average Compartment
1	141.23	2514	242.39	2897.62	106	24	2.208
2	146.43	2514	242.63	2903.06	106	24	2.292
3	135.72	2588	240.17	2963.89	106	23	2.217
4	146.27	2531	241.4	2918.67	106	23	2.391
5	143.83	2491	242.39	2877.22	106	23	2.348

106n_106o_7tu_3s (TestC.2) 50%

Seed	Loading cost	Traveling cost	Unloading cost	Total cost	Stops	Routes	Average Compartment
1	203.81	2476	362.48	3042.29	106	23	2.217
2	203.81	2503	362.11	3068.92	106	23	2.217
3	183.84	2557	359.16	3100	106	23	2.000
4	207.71	2525	362.11	3094.82	106	23	2.261
5	190.94	2510	354.36	3055.3	106	24	2.000

106n_106o_7tu_3s (TestC.2) 100%

Seed	Loading cost	Traveling cost	Unloading cost	Total cost	Stops	Routes	Average Compartment
1	233.79	2543	471.5	3248.29	106	24	1.833
2	255.84	2542	481.34	3279.18	106	24	2.000
3	244.81	2544	478.38	3267.19	106	23	2.000
4	196.45	2521	465.59	3183.04	106	23	1.609
5	226.24	2501	481.34	3208.58	106	23	2.174

References

- Chen, L., Liu, Y. and Langevin, A. (2019) 'A multi-compartment vehicle routing problem in cold-chain distribution', *Computers and Operations Research*, 111, pp. 58–66.
- Clarke, G. and Wright, J.W. (1964) 'Scheduling of vehicles from a central depot to a number of delivery points', *Operations Research*, 12(4), pp. 568–581.
- Coelho, L.C., Cordeau, J.-F. and Laporte, G. (2012) 'Consistency in multi-vehicle inventory-routing', *Transportation Research Part C: Emerging Technologies*, 24, pp. 270–287.
- Derigs, U., Gottlieb, J., Kalkoff, J., Piesche, M., Rothlauf, F. and Vogel, U. (2011) 'Vehicle routing with compartments: applications, modelling and heuristics', *OR Spectrum*, 33(4), pp. 885–914.
- Dueck, G. (1993) 'New optimization heuristics: the great deluge algorithm and the record-to-record travel', *Journal of Computational Physics*, 104(1), pp. 86–92.
- Eshtehadi, R., Demir, E. and Huang, Y. (2020) 'Solving the vehicle routing problem with multi-compartment vehicles for city logistics', *Computers and Operations Research*, 115, 104859.
- Fernie, J. and Sparks, L. (eds.) (2009) *Logistics and retail management: emerging issues and new challenges in the retail supply chain*. 3rd edn. London: Kogan Page.
- Frank, M., Ostermeier, M., Holzapfel, A., Hübner, A. and Kuhn, H. (2021) 'Optimizing routing and delivery patterns with multi-compartment vehicles', *European Journal of Operational Research*, 293(2), pp. 495–510.
- Glatzel, C., Großpietsch, J. and Hübner, A. (2012) 'Higher margins through efficient supply chains', *Akzente*, 2, pp. 16–21.
- Grünrock-Kern, U. (2013) *Deutscher Logistik-Preis 2013 für Lekkerland*. Bremen: Bundesvereinigung Logistik. <http://www.bvl.de/presse/meldungen/archiv-2013/deutscher-logistik-preis-2013-fuer-lekkerland>
- Gunawan, A., Kendall, G., McCollum, B., Seow, H.-V. and Lee, L.S. (2021) 'Vehicle routing: review of benchmark datasets', *Journal of the Operational Research Society*, 72(8), pp. 1794–1807.

- Hübner, A. and Ostermeier, M. (2019) 'A multi-compartment vehicle routing problem with loading and unloading costs', *Transportation Science*, 53(1), pp. 282–300.
- Kohan, E., Torres, F., Silva-Febré, V. and Pina-Pardo, J.C. (2026) 'Designing multi-compartment last-mile vehicle fleets: an open-source matheuristic', *Computers & Industrial Engineering*, 219, 112134.
- Kuhn, H. and Sternbeck, M.G. (2013) 'Integrative retail logistics: an exploratory study', *Operations Management Research*, 6(1), pp. 2–18.
- Laporte, G. (2009) 'Fifty years of vehicle routing', *Transportation Science*, 43(4), pp. 408–416.
- Ly, C. and Xu, J. (2023) 'Multi-compartment vehicle selection and delivery strategy with time windows under multi-objective optimization', *Alexandria Engineering Journal*, 85, pp. 146–159.
- Masmoudi, M.A., Baldacci, R., Fontaine, P. and Khayya, E. (2026) 'A novel multi-compartment vehicle routing problem with time windows: reconfigurable vehicles and dynamic compartments', *European Journal of Operational Research*, in press.
- Muyldermans, L. and Pang, G. (2010) 'On the benefits of co-collection: experiments with a multi-compartment vehicle routing algorithm', *European Journal of Operational Research*, 206(1), pp. 93–103.
- Ostermeier, M. (2024) 'The supply of convenience stores: challenges of short-distance routing within the constraints of working time regulations', *European Journal of Operational Research*, 314(3), pp. 997–1012.
- Ostermeier, M. and Hübner, A. (2018) 'Vehicle selection for a multi-compartment vehicle routing problem', *European Journal of Operational Research*, 269(2), pp. 682–694.
- Ostermeier, M., Henke, T., Hübner, A. and Wäscher, G. (2021) 'Multi-compartment vehicle routing problems: state-of-the-art, modeling framework and future directions', *European Journal of Operational Research*, 292(3), pp. 799–817.
- Ropke, S. and Pisinger, D. (2006) 'An adaptive large neighborhood search heuristic for the pickup and delivery problem with time windows', *Transportation Science*, 40(4), pp. 455–472.

Shaw, P. (1997) *A new local search algorithm providing high quality solutions to vehicle routing problems*. Technical report. Glasgow: Department of Computer Science, University of Strathclyde.