MASTER THESIS

DESIGNING THE LOGISTICS NETWORK FOR A PARCEL CARRIER

J.P.R. VAN PIJKEREN

UNIVERSITY OF TWENTE.

Company X

UNIVERSITY OF TWENTE
FACULTY OF BEHAVIOURAL, MANAGEMENT AND SOCIAL SCIENCES

DESIGNING THE LOGISTICS NETWORK FOR A PARCEL CARRIER

MASTER THESIS
INDUSTRIAL ENGINEERING & MANAGEMENT
PRODUCTION AND LOGISTICS MANAGEMENT

This is a publicly available version of this thesis. Therefore, all confidential datapoints are multiplied by a random number such that they are not representative of reality. Furthermore, all figures and tables with non-crucial sensitive information will be shielded.

AUTHOR

J.P.R. van Pijkeren

DATE

31-03-2023

EXAMINATION COMMITTEE
University of Twente

Dr. ir. E.A. Lalla-Ruiz
Assistant Professor in Logistics and Operations Research
Dep. Industrial Engineering and Business Information Systems
Faculty of Behavioural Management and Social Sciences

Dr. A. Trivella
Assistant Professor of Operations Research
Dep. Industrial Engineering and Business Information Systems
Faculty of Behavioural Management and Social Sciences

External committee member

although
MANAGEMENT SUMMARY

Context

This research is conducted at a company that wants to remain anonymous, it is referred to as Company X. Company X is a start-up parcel carrier in the Netherlands founded at the end of 2021. Company X was founded by a B2C e-commerce company (ECC) which sells products from its own stock as well as products from other large and small retailers all over the Netherlands. Since these products need to be shipped from these retailers to the consumers' homes, ECC arranges the logistical fulfilment of collecting and delivering these products. It is the strategic ambition of the ECC to have full control over the logistical fulfilment of its own parcels and the parcels from other retailers that are sold through the ECC.

The logistical network is the foundation of Company X's operations and is therefore of great importance to achieve its ambition of being the most reliable, efficient and sustainable parcel carrier. Company X is scaling fast and one of the challenges it is facing is the design of its logistical network. Currently, Company X has an initial plan made for them on how to design its network. However, this approach is based on simplifications and assumptions and does not fully enable Company X to achieve its ambition. Furthermore, with a changing market, Company X wants to revise this initial plan, for which it does not have the right tooling.

To be able to achieve its ambition, Company X wants to become able to periodically determine the best network design to achieve the most efficiency in its operations. The best network minimizes the total operational costs while also having a robust performance over multiple scenarios. Operational costs include:

- The operational costs of opened depots (parcel sorting, cost of materials).
- The costs of inter-hub transportation (linehaul).
- The costs involved with the routes that are driven from each depot to the retailers.

Therefore, the goal of this research is to build a model that will enable Company X to design its network such that it can operate at the lowest operational costs and periodically analyse its network performance. The research question of this thesis is formulated as:

What is the best location of and allocation to depots such that the logistical network of Company X operates at minimal costs?

Method

The main elements that need to be decided on in the design of the logistical network are the locations of hubs, linehaul, and which locations to service from which hub. From the literature, it becomes clear that the network of Company X can be modelled as a hub-and-spoke (HS) network. An HS network is characterized as a network type, used in various many-to-many systems, with hubs and non-hubs where large volumes are transported between hubs to achieve economies of scale and non-hubs are used to distribute volumes, passengers or information over the network.

The theoretical problem that describes the challenge of designing the logistical network best is described in the literature as the 'hub location routing problem' (HLRP). HLRPs are concerned with the optimization of an HS network by integrating the hub locations, inter-hub structure, and operational routing cost. Using this literature, we have formulated the design of the logistics network of Company X as a mixed-integer linear program. However, due to the size of the problem, the Netherlands consists of 4070 postal zones, it is intractable to solve this model to optimality. Therefore, we propose a heuristic approach based on a general variable neighbourhood search (GVNS) to find the best possible network within a reasonable time. This GVNS approach is a combination of elements from the fastest heuristic approach found in literature and the heuristic approach that has previously been able to solve a problem in a similar (although much smaller) context.

A disadvantage of this GVNS method is that it is a deterministic local search and is therefore not able to grasp the stochasticity and fluctuations in workload that Company X must handle. Therefore, we have expanded the

deterministic heuristic with a simulation component into a simheuristic-based GVNS. In the simheuristic, all promising solutions that are found while solving the deterministic GVNS are moved to a simulation procedure. This procedure evaluates the solution under multiple scenarios so that the best solutions are those who perform best under different circumstances. Furthermore, the best-performing promising solutions are put into an even more intensive simulation to gain insights into their stochastic behaviour as thoroughly as possible.

One of the most important inputs to the model is the set of considered hub locations. To limit the required computation time, we cannot put in many hub locations so that the solution space of the model does not become too large. Therefore, based on a dataset of all industrial/commercial areas in the Netherlands, we have determined a limited set of 30 possible hub locations. In this research, hub locations and service areas are determined on a PC4 postal zone level. A PC4 postal zone contains the area in which the streets lie with the same four digits in the postal code. The Netherlands is divided into a total of 4070 PC4 postal zones.

Results

Using historical sales of the ECC, we have performed three different experiments to find the best possible network for Company X. In the first we used our proposed simheuristic-based GVNS to determine the best network for Company X. Secondly, we tested the GVNS but with a much faster simple allocation strategy for postal zones. Lastly, we assessed the performance of the network when we take the hub locations from the initial plan as input.

Using the simheuristic-based GVNS and the set of possible hub locations, we have found a logistics network that has an average daily cost of operating of €25,290. Using the simple allocation strategy, we found a slightly better solution with an average daily operating cost of €24,739. Lastly, when taking the initial plan as input, the best solution we found has a daily operating cost of €27,262, which is worse than the other methods.

Conclusions and recommendations

In this research, we have created a model that describes the logistical network of Company X. The model enables Company X to analyse its network and build it such that it can operate at minimal costs. By expanding the GVNS with a simulation component into a simheuristic, we enable Company X to not only solve deterministic scenarios but evaluate the stochastic behaviour of these solutions. Furthermore, we performed three experiments in which we used our model in different ways to find the best network. We found a network design that improves the initial plan.

While these results are positive, we have some recommendations for using the model as well as improving it. Because the GVNS searches for solutions in a deterministic scenario, the solutions are built based on that scenario. Using another deterministic scenario might lead to different choices for hub locations or service areas. Therefore, we advise Company X to solve the model for multiple different deterministic scenarios and make a final decision for the network design based on the simulated results of these solutions. In this way, the influence of a single scenario on the local search is decreased. Additionally, we have proposed multiple improvements to Company X and a direction for future work. These include, among others, improving the hub level operators and investigating the added value of using a multi-allocation HLRP formulation.

PREFACE

Dear reader,

Before you lies my thesis named 'Designing the best logistical network for a parcel carrier'. With the completion of this thesis not only a challenging 10-month project but also my time as a student has come to an end. I have had the pleasure to spend six years of my student life in Enschede, having followed both the bachelor and master programmes of Industrial Engineering & Management. I can look back on a wonderful period in which I learned a lot both curricular and extracurricular through two years in the faculty council, many committees and a board year at Study Association Stress.

I want to thank Eduardo Lalla-Ruiz and Alessio Trivella for their work as supervisors at the University of Twente. Their critical feedback has helped me in improving the academic quality of my thesis. Even though we had some discussions on the contents, they have always been very supportive and helped to make sure I can finish my thesis on time.

I also want to thank the team with whom I have had the pleasure to spend the last 10 months with on a daily basis. They have made this assignment possible, have been very supportive and helped me make sure I could add value to the organisation. A special thanks to Jan, who has put in tremendous effort in helping me. From first having weekly meetings to daily check-ins, Jan has really thought along during the whole process and helped me both in formulating the model as well as scoping. This thesis would not have come together without his help.

Furthermore, I want to thank my friends, family and girlfriend for making these last six years as fun as they have been and for supporting me during this thesis project.

I hope you enjoy reading this thesis!

Jelle van Pijkeren
Utrecht, March 2023

TABLE OF CONTENTS

Management Summary iv
Preface vi
Table of Contents vii
List of Figures ix
List of Tables x
List of Algorithms xi
Glossary of terms and abbreviations xii
1 Introduction & Research Design 1
1.1 Research Context 1
1.2 Problem Analysis 1
1.3 Research Design 4
2 Context Analysis 7
2.1 Postal zones in the Netherlands 7
2.2 Market Analysis 7
2.3 Logistical Processes 11
2.4 Network design for Company X 14
2.5 Transport Resources 15
2.6 Conclusions on the context analysis 16
3 Literature Review 18
3.1 Network Design Problems 18
3.2 Hub-and-Spoke Networks 20
3.3 Hub Location Problem 21
3.4 Hub Location Routing Problem 23
3.5 Solution Approaches 30
3.6 Decision-making in real-life applications 32
3.7 Conclusions on literature 36
4 Modelling the optimal network 37
4.1 Assumptions and approximation on postal zones 37
4.2 HLRP definition of the network of Company X 37
4.3 Mixed-Integer Linear Programming Model 40
4.4 GVNS-based heuristic 46
4.5 Simheuristic-based GVNS 51
4.6 Conclusions on modelling the best network 52
5 Numerical Experiments 53
5.1 Experimental Design 53
5.2 Data Instances 54

5.3 Model tuning, testing, and validation 56
5.4 Results of simheuristic-based GVNS 63
5.5 Results compared with other methods 63
5.6 Sensitivity Analysis 65
5.7 Conclusion on numerical experiments 67
6 Conclusions & Recommendations 68
6.1 Conclusions 68
6.2 Recommendations for future research and limitations 68
6.3 Contributions to literature and practice 69
Bibliography 71
Appendix A. Results of experiment 1 75
Appendix B. Results of experiment 2 77
Appendix C. Results of experiment 3 80
Appendix D. Results of experiments 6, 7, and 8 83
Appendix E. Results of experiment 8 84
Appendix F. Results of experiment 9 87

LIST OF FIGURES

Figure 1 The ambition and challenge of Company X 2
Figure 2 Total revenue of e-commerce sales in the Netherlands from 2005 to 2021 (Statista, 2022a) 8
Figure 3 Total number of retailers PC4 area 9
Figure 4 Development of unique retailers per week 9
Figure 5 Average demand per weekday (0 is Monday) over 2.5 years. 10
Figure 6 Difference in volume per day over the course of 2021 10
Figure 7 The distribution of average PPC volumes on Monday and the corresponding destinations (in parcels) 11
Figure 8 Schematic graph of PPC network 12
Figure 9 Schematic graph of BPC network 12
Figure 10 Overview of different processes over the day (for abbreviations, see Table 1 and Table 2) 13
Figure 11 Candidate hub locations 16
Figure 12 An overview of how different network design elements are covered in theoretical known problems. 18
Figure 13 Structure of (A) Cycle-star, (B) Star-star, (C) tree-star, and (D) line hub networks (Contreras, 2021) 23
Figure 14 Overview of solution methods for the Hub Location Routing Problem 30
Figure 15 General scheme of simheuristics for solving stochastic COPs (Juan et al., 2015) 35
Figure 16 Example of G, without edges 38
Figure 17 An example of subgraph C_1, the PPC process 39
Figure 18 An example of subgraph C_2, the BPC process 39
Figure 19 An example of subgraph C_3, the HD process 39
Figure 20 An example of the resulting linehaul of the subgraphs C_1 and C_3 40
Figure 21 Solution of the MILP model for a very small test instance of 10 nodes 45
Figure 22 Graphical representation of working of GVNS algorithm on multiple levels. 50
Figure 23 Confidence interval half-width relative to the mean for $\alpha = 0.05$ 61
Figure 24 Confidence interval half-width relative to the mean for $\alpha = 0.02$ 62
Figure 25 Confidence interval half-width relative to the mean for $\alpha = 0.01$ 62
Figure 26 Network design determined by the simheuristic-based GVNS. 63
Figure 27 Network design using the simple allocation method. 64

LIST OF TABLES

Table 1 An overview of the used definitions 13
Table 2 An overview of the different logistical processes 13
Table 3 Vehicle characteristics 16
Table 4 Abbreviations not earlier introduced 28
Table 5 An overview of literature on the Hub Location Routing Problem 29
Table 6 Results of experiment 1: Parameter tuning 56
Table 7 Results of experiment 2: Operator performance 57
Table 8 Results of experiment 2: Operator order 58
Table 9 Results of experiment 3: Impact of the initial solution 59
Table 10 Results of experiment 4: Model validation 60
Table 11 Results of the simple approach compared to simheuristic-based GVNS for test instances 64
Table 12 Results of experiment 9: Sensitivity Analysis 66
Table 13 Results of experiment 1 75
Table 14 Results of experiment 2: Operator performance 77
Table 15 Results of experiment 2: Operator order 78
Table 16 Results of experiment 3 80
Table 17 Results of experiments 6, 7, 8 83
Table 18 Results of experiment 9: Sensitivity analysis - objective 87
Table 19 Results of experiment 9: Sensitivity analysis – opened hubs 87
Table 20 Results of experiment 9: Sensitivity analysis –Hub opening costs 87
Table 21 Results of experiment 9: Sensitivity analysis – Linehaul costs 88
Table 22 Results of experiment 9: Sensitivity analysis – Total routing costs 88
Table 23 Results of experiment 9: Sensitivity analysis – Total number of routes 88
Table 24 Results of experiment 9: Sensitivity analysis – total route length 89
Table 25 Results of experiment 9: Sensitivity analysis – routing duration 89

LIST OF ALGORITHMS

Algorithm 1: Initial Solution 47
Algorithm 2: General Variable Neighbourhood Search 49
Algorithm 3: Shaking Procedure 50
Algorithm 4: Multi-Level Variable Neighbourhood Descent 51
Algorithm 5: General Variable Neighbourhood Search with simulation 52
Algorithm 6: Simulation Procedure 52
Algorithm 7: Creation of simulation scenarios 56

GLOSSARY OF TERMS AND ABBREVIATIONS

Term/Abbreviation	**Definition**
ECC	E-commerce Company
ECW	E-commerce Company’s Warehouse
PP	Piecewise Parcels
BP	Batch Parcels
PPC	Piecewise Parcel Collection
BPC	Batch Parcel Collection
HD	Home Delivery
LH	Linehaul
FS	First Sort
SS	Second Sort
PuR	Pickup Retailer
DR	Deposit Retailer
BR	Batch Retailer
CP	Collection Point
HS	Hub-and-Spoke
HLP	Hub Location Problem
LRP	Location Routing Problem
HLRP	Hub Location Routing Problem
GVNS	General Variable Neighbourhood Search
VND	Variable Neighbourhood Descent
CVRP	Capacitated Vehicle Routing Problem
KPI	Key Performance Indicator
PC4	Postal Code with 4 digits
LTL	Less-Than-Truckload
Hub/Depot	A consolidation centre where routes depart and arrive and where parcels are sorted.
Service Area	The set of locations that together form the stops on the routes that are driven from one depot.
Linehaul	The transportation of consolidated parcels from the hub where they are collected to the hub from where they should be delivered.

1 INTRODUCTION & RESEARCH DESIGN

In this chapter, we introduce the companies that are involved in this research (Section 1.1). Besides this, the chapter provides information about the problem analysis and the scope of the research (Section 1.2), and the research design (Section 1.3). In Chapter 2, the relevant processes and the problem are explained.

1.1 RESEARCH CONTEXT

In this section, we introduce the companies that are involved in this research the E-commerce Company and Company X, where the research is conducted, are introduced.

1.1.1 E-commerce Company

The second of the three stakeholder companies is the e-commerce company (ECC). The e-commerce company is a Business to Consumer (B2C) company that sells all kinds of consumer products from its warehouse. Next to their own products, ECC also sells products that come from other retailers. The logistical fulfilment for these shipments is arranged by ECC. So, ECC contracts carriers to perform the logistics for the retailers. It is the strategic ambition of the ECC to have full control over the logistical fulfilment of its own parcels as well as the parcels from other retailers that are sold by the ECC. To achieve that, Company X is built.

These products are shipped directly from the retailer to the consumer. These products must be collected first before they can be delivered to consumers' houses. For other retailers, ECC makes sure that large batches of a product are brought to their own warehouse from where they are shipped by parcel carriers. Therefore, the logistical services that are performed can be separated into different groups based on the way ECC handles the products from third-party retailers. First, there are the products that are collected piecewise (piecewise parcels, PP) and have to be sorted to hand them over to other parcel carriers. Second, some products are shipped to the ECC's warehouse in large batches (batch parcels, BP). Both these product groups must be collected and sorted by Company X.

1.1.2 Company X

Thirdly, Company X is the company at which this research is conducted. Company X is a start-up that was founded at the end of 2021 and provides logistical services on behalf of the B2C E-commerce Company. They are responsible for the transportation of parcels. Although Company X is an independent company, it was created by and works solely for ECC. Company X was created to fulfil the strategic ambition of the ECC to have all parcel transhipments done by themselves and thus not have other third-party parcel carriers have their products shipped to consumers.

Company X wants to be a disrupting new player in the market of logistical fulfilment. Therefore, it is its ambition to offer the best parcel transportation service possible by making sure it is:

- The most **reliable**, meaning Company X can deliver what they promise.
- The most **efficient**, offering it against the most competitive price.
- The most **sustainable**, by making sure it is scalable and the environmental impact is as low as possible.

At this point, Company X is only performing a small part of the collection of the two parcel types as described above (PP/BP). The rest is performed by other carriers. It is their ambition to grow fast and take over at least half of all shipments. Furthermore, Company X wants to investigate whether they can expand their services by also doing home delivery.

1.2 PROBLEM ANALYSIS

In this section, the problem is analysed in more detail. We address the challenges and desires of Company X that have led to the formulation of the core problem of this thesis. Furthermore, we elaborate on the scope of the research. For an explanation of the different logistical processes that are relevant to this research, see Section 2.3.

1.2.1 Problem Context

Although Company X has only been founded at the end of 2021, it is growing rapidly due to the support of the ECC. Although many processes within the company have been worked out and set up, there remain multiple decisions and processes that must be worked out in detail. This includes the design of the logistical network, which is an important aspect of the operations of a parcel carrier as it sets the conditions in which the daily operations can be performed, and thus is very important for Company X to achieve their ambition. Within this network, multiple depots need to be placed from where routes are driven to collect (or deliver) parcels.

To further explain the challenges that Company X is facing, the definition of service area must be given. A service area of a depot is the set of locations that together form the stops on the routes that are driven from that depot. As also further explained in Section 2.3.1, these stops can also include collection points. Deposit retailers (DR), which are retailers with small sales quantities, must bring their parcels to a collection point. The allocation of DRs to collection points is thus also part of the service area.

Since Company X has just been founded, the logistical network, which includes depot locations and which areas to service from which depot, is yet to be built. Company X has an initial plan made for them on where to open depots and corresponding service areas. However, this plan is based on many assumptions and approximations and therefore does not fully enable Company X to fulfil its ambition. Therefore, Company X wants the network design to be more advanced and improved by means of an optimization/improvement algorithm. Because Company X wants to be the most reliable, efficient, and sustainable parcel carrier in the market, it wants to make sure that it builds its network such that it can operate most efficiently in the short as well as the long term. With the operation the following costs are included:

- The operational costs of opened depots.
- The costs of inter-hub transportation (linehaul).
- The costs involved with the routes that are driven from each depot.

Therefore, being able to analyse and find the best network design based on the costs above and given any forecast enables Company X to always have the blueprint for an optimized network when opening new depots. This can save costs both operating costs as well as changeover costs.

Furthermore, there exist great seasonal effects on the workload for company X within both weeks as well as months. First of all, there is a peak workload on Mondays or Tuesdays, depending on the logistical process (see Section 2.2.3). The volumes on the other three days of the week are always lower. Furthermore, there exists great seasonality within a year, the first three quarters of the year are, overall, quite similar and steady in terms of volumes per week. Yet in the last quarter of the year, there is a big increase in the volume. The volume quantities are (more than) four times the volume per week compared to the other months of the year. Dealing with this peak(s) in managing their logistical network is a great challenge for Company X.

On top of that, Company X currently relies heavily on the data and forecasts that are provided to them by ECC. The quality of this data cannot be verified other than by comparing it to what happens during daily operations. Besides this, the data is very aggregated, meaning it is hard to translate the forecast into workload and routes for specific service areas.

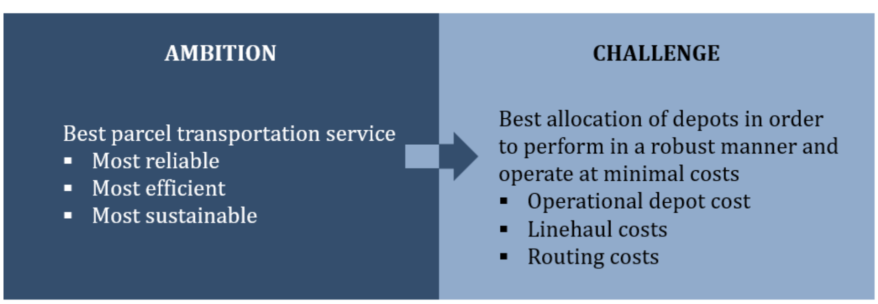

Figure 1 The ambition and challenge of Company X

Therefore, to be able to operate more independently from ECC, Company X wants to build the capability to be able to assess and analyse strategic and tactical scenarios with a tool. A tool with these capabilities will enable Company X to analyse the impact of changes on their network and enables them to decide which actions need to be taken under which scenario. Types of questions that can be thought of in this regard are 'how many extra vehicles do we need of the workload in area X increases by Y?'.

1.2.2 Core Problem

As described, Company X is going to build its network and thus must decide on depot locations and corresponding service areas. To achieve its ambition, Company X wants to find the network that minimizes operational costs while also having a robust performance. Besides this, Company X has the desire to be able to make scenario analyses on both the strategic and tactical levels. With the dependency on the forecast of the ECC and the great seasonal peaks, Company X must be able to independently determine what the best course of action is given a certain scenario or how it can organise itself as efficiently as possible.

Thus, the challenge is to determine the best logistical network for Company X given a certain scenario. To solve this challenge, multiple smaller elements must be determined that together shape how the logistical network should look like, this includes:

1. The location of depots, where collected parcels are sorted and routes depart and arrive.
2. The sizes of the depots.
3. The service area of the depots.
4. The location of the collection points.
5. The retailers that are connected to a collection point.
6. The linehaul structure between depots.
7. Which routes will be used.

When it is possible to find the best network for a given scenario it becomes possible to determine the best network given a set of scenarios. Thus, providing the best network for Company X.

1.2.3 Scope of the research

As described in section 1.2.2, the core problem, determining the best logistical network for Company X, consists of seven elements. Since this research should be conducted within a time frame, we must pay attention to the scope of the research. Therefore, we have decided to not take all seven logistical challenges into account. However, we want to aim for the broadest possible approach and we, therefore, proceed as follows.

The main goal of this research, *determining the best logistical network*, is a strategic question. This question contains the location and size of the depots and the corresponding service areas. So, the first three elements should be taken into account in this research.

At first sight, linehaul and routes (questions six and seven) do not seem to be questions to be included in this strategic research since these are more operational questions. However, it is important to take them into consideration when determining the location and size of the depots and the service areas. Since the goal is to figure out the most cost-effective network and this can be done from scratch with no changeover costs, it is important to not only take into account the locations of depots, but also the costs of the whole network, including costs of linehaul and routing. Thus, considering both linehaul and routing provides the most realistic picture in terms of costs, and therefore the opportunity to build the best network. One can agree that routing is an operational decision and will vary every day, but for Company X in practice, routes are often similar every day. This makes sense since the pickup retailers are large retailers, so every day they have volume to be shipped. And therefore, the locations that need to be visited are (almost) always the same. The same goes for delivery, postal zones with a high population density are also the location where a lot of packages need to be delivered.

Thus, the questions about depots, service areas, linehaul, and routes form the foundation of the logistical network and therefore we think it is important to first answer these. To keep the size of the research manageable, we decided to leave out the fourth and fifth elements, the location of the collection points, and the connection between retailers and the collection points.

The reason for this is that when considering where to locate the collection points, every retailer can become a collection point, which broadens the solution space tremendously. However, it is possible to take this into account in the model by predetermining it before solving the model and thus not considering it during the improvement phase. So, only making the decision once. Based on data and with the help of a greedy heuristic, we can estimate locations for collection points and the connection between retailers and the collection points. In this way, we can select the best network while taking into account all aspects, and declining the solution space with this delineator. The explanation of how the locations of collection points are determined and which retailer is connected to which collection point is explained in Section 4.4.1.

To conclude, we decided that, to determine the best logistical network for Company X, we should take into account decisions on depots, service areas, linehaul, and routing. We also decided to predetermine the location of collection points and the connection between retailers and collection points to keep the problem and therefore solution space manageable in the limited time we have for this research.

1.3 RESEARCH DESIGN

In this section, the research design is described. The research goal is explained first. Whereafter, we discuss the research questions.

1.3.1 Research Goal

As derived from the problem analysis, the challenge of this research is finding the best logistical network for Company X. This comprises finding the best locations and sizes for depots as well as determining the corresponding service areas to each depot such that the expected routing costs are minimal while having nationwide coverage in the Netherlands. This challenge covers both the collection of parcels from retailers and the home delivery to consumers. As mentioned before, the logistical network is not yet built.

The desired outcome of this research should be an approach (algorithm) that is able to use input like forecast, potential hub locations, service levels, and routing costs to determine the best logistical network for Company X. Next to that, Company X wants to create the tool to do strategic and tactical analyses. Therefore, the tool needs to be practical enough to be reusable, e.g., by analysing a forecast and determining if extra resources are needed given a set of available depots.

These outcomes should be realized by making an (optimisation) algorithm that combines locating the depots, sizing the depots, and dedication of service areas such that the cost of driving the routes to visit these locations from the depots and the depot costs is minimized. The precise approach that the algorithm will follow will be based on the literature review.

1.3.2 Research Questions

As explained in the problem analysis, the goal of this thesis is to determine what the logistical network of Company X should look like. Therefore, the main question of this research is:

What is the best location of and allocation to depots such that the logistical network of Company X operates at minimal costs?

By answering the following sub-questions, we answer the main research question:

1. ***How is the current logistical network design done for Company X?***

Before we start with researching and creating models, we first have to create a good understanding of the work that Company X has to do, which logistical processes have to be performed through the network, how the current network has been designed, and which resources are available. By knowing this, we know what has to be covered by the model we will make. This will also form the input for the mathematical model. Here we can think of the number of parcels that need to be collected, but also the indirect inputs like how many parcels fit in a vehicle. This means that based on historical data want to determine:

- How is the workload for Company X divided over the Netherlands?
  - How is the Dutch e-commerce market evolving?
  - The division of collection (and delivery) volumes over the country.
  - Seasonal effects within weeks and within a year.
  - Locations, volumes, and service times of retailers.
- What are the logistical processes that need to be performed through the network?
- How is the current logistical network designed?
  - What are the requirements?
  - What are the Key Performance Indicators?
- What will be the available set of resources?
  - The set of available vehicles (capacity, operating cost etc.).
  - The set of depot types (capacity, operating cost, etc.)

We elaborate further on these questions in Chapter 2.

2. ***Which models exist in the literature that optimize transportation networks with depots and hubs?***

The first step for answering the main research question is to review the status quo in the literature regarding the design of logistical networks. We will search for sources that describe models that solve the same type or similar types of logistical networks. This includes:

- What are the mathematical formulations of such models?
- What are the (heuristic) approaches they use to solve these models?
- What are the computational times of these approaches?
- What are the strengths and weaknesses of these approaches?
- Are there similar cases known?

With the combined literature, we know whether there are proven methods to model the logistical network of Company X or whether we need to make adjustments or additions to these models.

The answers to these questions can be found in Chapter 3.

3. ***How can we model and solve the logistics network of Company X?***

From the literature review that is created by answering research question two, we know to which extent we can use known methods to model the logistical network. From there, we can make a mathematical model for Company X, if needed by making adjustments to known methods. Subsequently, based on the literature, we will design a (heuristic) approach to solve the mathematical model efficiently. To make a good model, we need to find answers to the following questions:

- What are the assumptions needed to model the network?
- What is the scope of the model?
- How can we solve the logistical network of Company X?

The model is explained in Chapter 4.

4. ***What is the best logistics network for Company X in the Netherlands?***
   a. ***What are the best parameters for the algorithm, and can we validate the outcomes?***

When the input analyses, mathematical model, and algorithm stand. It is time to test the model. Firstly, the right model parameters must be found, this means the parameters that give the best objective values within reasonable timeframes. Furthermore, we need to make sure that the outcomes of the model are an accurate reflection of reality and thus can be trusted.

b. What is the best network for Company X?

After we have determined the best model parameters and validated that the model results are accurate representations of reality. We want to let the model find the best logistical network for Company X.

c. How does this network perform compared to the initial plan and simple methods?

Having found a network design using our model, we will test whether the model is actually of added benefit for Company X. Firstly, this will be done by comparing the results to the outcomes a simple approach that takes less time to solve. Secondly, by comparing the results to the depot locations as they are chosen in the initial plan of Company X.

d. What is the sensitivity of input parameters?

Lastly, we want to know what the impact of certain input parameters, like depot costs or vehicle capacities, is on the outcome of the model.

The experiments and corresponding results can be found in Chapter 5.

5. *What are the conclusions of the research and recommendations for Company X?*

Lastly, when all the other research questions have been answered the findings have to be put together. In this section we will formulate a conclusion to the main research question, thus showing the optimal logistical network. Furthermore, recommendations will be formulated for Company X on how to use this research in the future so that they can make full use of it.

The conclusion and recommendations can be found in Chapter 6.

2 CONTEXT ANALYSIS

In this chapter, we investigate the context in which our problem exists so that the research question '*How is the current logistical network design done for Company X?*' can be answered. We start by giving an overview of the (expected) workload for Company X (Section 2.2). This includes general trends in the Dutch e-commerce market and the position of ECC in it. Furthermore, we present an analysis of historical sales data of the past 2.5 years from which we show the following characteristics for future 'demand' of Company X:

- The division of collection and delivery volumes over the country.
- Seasonal effects between days of the week and within a year.
- The set of available vehicles (capacity, operating cost etc.).
- Locations and volumes retailers.

Next, we present an overview of the logistical processes that are performed by Company X (Section 2.3), explain the current network design process and the used key performance indicators (Section 2.4), and give an overview of the vehicle types and depot characteristics that will be included in the network design (Section 2.5).

2.1 POSTAL ZONES IN THE NETHERLANDS

Within the Netherlands, every street has a postal code. This postal code consists of four digits and two letters, which together form a six-element ID, e.g., 1011 AA, that is called the PC6. Each PC6 corresponds to (parts) of a street such that each PC6 has a limited number of houses/buildings related to it. The next level is the PC4, which consists only of the digits in the postal code. The PC4 will be the main level that is looked at in the research, therefore when the term 'postal zone' is used the PC4 level is meant. The PC4 level spans over neighbourhoods and thus contains multiple streets. The area size of a PC4 area differs strongly, typically the area is bigger in more rural areas. Each PC4 is part of a PC3 area, which is the area of all PC4 areas that share the same first three digits, e.g. all postal zones with 145_ __, these areas lie next to each other and form a cluster. The last level that is relevant for this research is the PC2 level, which is the two-digit equivalent of the PC3 level.

2.2 MARKET ANALYSIS

In this section, we give an overview of the e-commerce market in the Netherlands, touch upon the market position of the e-commerce company and translate that into the workload for Company X in terms of the number of stops, volumes to collect, seasonal effects and how the volume is distributed over the country.

2.2.1 E-commerce Market in the Netherlands

The one customer of Company X is the e-commerce company (ECC) (see Section 1.1.1). This means that although Company X is not a player in the e-commerce market the developments of this market are of great importance to the business of Company X.

2.2.1.1 Market size

The e-commerce market is big in the Netherlands, with little over 30.5 billion euros it occupies about 84% of the digital revenue market. Moreover, in 2021 86.7% of all Dutch individuals shop online (Statista, 2022a), which is higher than the average of northern, central, and western Europe (Statista, 2022b). The market size of e-commerce in the Netherlands in terms of revenue has also been steadily increasing over the period 2005-2021 (see Figure 2) and is expected to do so with 12.5% per year from 2021 to 2025 (Statista, 2022b). The number of orders has grown with a relatively similar pattern as the revenue to a total of 373.4 million orders in 2021 (Statista, 2022a).

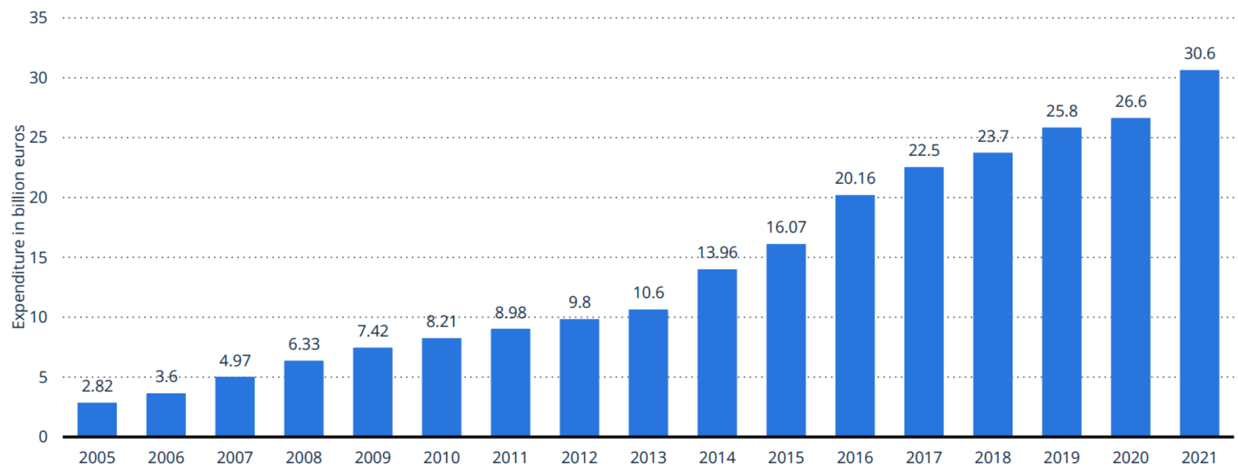

Figure 2 Total revenue of e-commerce sales in the Netherlands from 2005 to 2021 (Statista, 2022a)

Although there is a clear significant growth trend over the last few years, this growth has been decreasing strongly over the last months of 2022. Due to increased energy costs and inflation that is (expected to be) 3.4% higher than initially forecasted (Statista, 2022b) the sales in the e-commerce market have dropped strongly. This is also visible in the change of revenue rates of Q1 and Q2 of 2022 of -27% and -12% respectively compared to the same periods of 2021 (Statista, 2022a). It must be noted that these large decreases in revenue can also be explained by the country not being in lock-down anymore which can lead to different spending behaviour by consumers.

2.2.1.2 Competitive landscape

Just as the revenue, the number of online shops and mail order companies has increased (even faster) over the last years from 40,785 in 2019 to a total of 80,770 in 2022 (Statista, 2022a). In 2020, the five largest players in the market account for 37% of the total net sales, the top 6-25 for 29%, the top 26-100 for 23%, and the remainder 12% (Statista, 2022b). This shows that there are a few well-established big players followed by a small group of still relatively large companies, i.e., 100 companies account for 88% of the sales and thus 80,670 companies account for (only) 12% (assuming the division of sales in 2022 is equal to that of 2020). This shows that there are a lot of small companies/retailers in the market, of which multiple sell via ECC.

2.2.1.3 Consumer preferences

Lastly, research has been done into what the most important attributes are for consumers when doing online shopping. The most important attribute was a 'fast/reliable delivery', which relates directly to the service that Company X is offering, as mentioned by 40% of the almost 8800 respondents (Statista, 2022a). From this, we can see that having a good logistical network is of great strategic importance to e-commerce retailers, and thus to Company X. Furthermore, Dutch consumers are quite demanding in terms of the delivery time they expect. Around 30% of consumers expect a delivery time of 1-2 days, which is the highest of twelve European countries (Postnord, 2021). This implies that logistical networks transporting e-commerce sales in the Netherlands are under more pressure to deliver high service.

2.2.2 Market position e-commerce company

The e-commerce company (ECC) is a somewhat mature player in the Dutch e-commerce market[1]. They sell products in different types of product categories. Furthermore, they have been growing steadily, just above the market average. Therefore, we can assume that the growth rate of ECC will be equal to that of the whole market, i.e., 12.5% per year. Furthermore, ECC finds itself somewhere in the top 100 of e-commerce companies in the Netherlands, meaning their sales volume is around 80 million euros.

[1] *Due to confidentiality, we cannot disclose the exact numbers.*

2.2.3 Workload for company X

In this section, we present the analysis of historical sales data of 2.5 years that gives an overview of the factors that need to be considered when designing the logistical network. This is especially relevant for the input datasets that will go into the model.

2.2.3.1 Retailers

To determine the number of locations to visit and the location of collection points we need to know how many retailers there are, what their respective sizes are in terms of volume per day, what their locations are (Figure 3)[2], and how the number of retailers evolves over time (Figure 4)[2].

The number of retailers has increased quite a lot over the past two and a half years. Starting at selling products from only a few retailers a week in 2020 to well over 1300 retailers per week in 2022. In total, EEC has sold products from well over 3,455 unique retailers. Out of these retailers, only 200 meet the criteria to become a pickup retailer (PuR).

Figure 3 Total number of retailers PC4 area[2]

Furthermore, from the data, we know how these retailers are spread over the country. As can be seen from Figure 3, there are some clear hotspots with regard to where retailers are located. These are the Amsterdam and Rotterdam areas in the west, Groningen in the north, and some places in the east and south. All white areas are areas where no retailers are located.

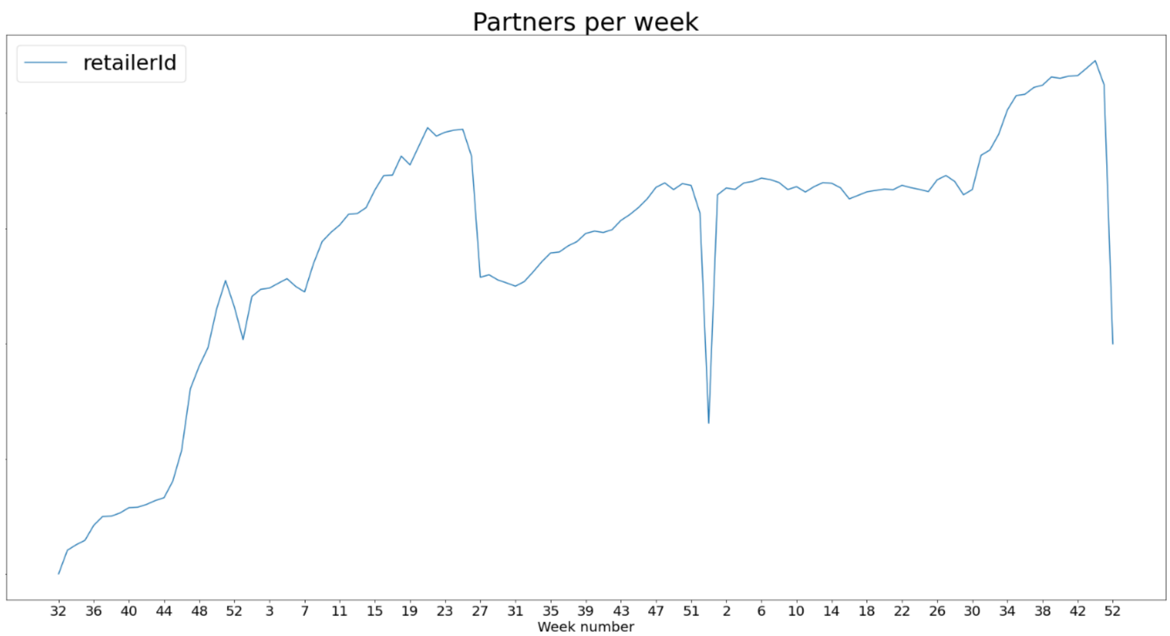

Figure 4 Development of unique retailers per week[2]

[2] *The y-axis has been removed to maintain confidentiality*

2.2.3.2 Volumes per day

In terms of the average volumes per day (see Figure 5), some observations can be made. First of all, there is a very clear weekly pattern to see in the data. Monday is the peak day, which follows logically from orders being placed over multiple days in the weekend. The Tuesday to Thursday are less but often quite the same. Lastly, a strong decrease can be seen on Friday and even stronger on Saturday. The total volume of packages that are sold through ECC can, at peaks, get over 15,000 items in a day (see Figure 6). In the short term, this volume will not completely be processed by Company X, but each year a larger part will be taken on.

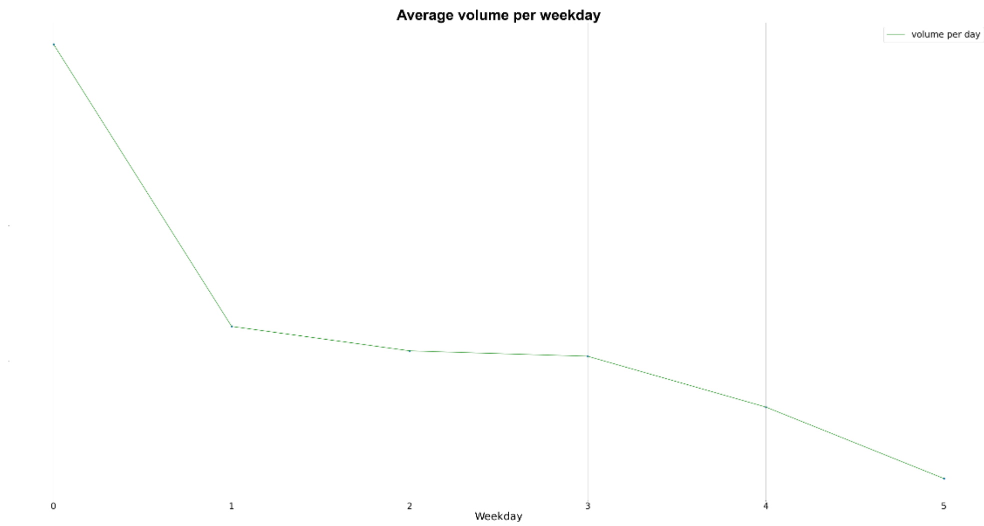

Figure 5 Average demand per weekday (0 is Monday) over 2.5 years[2].

2.2.3.3 Seasonal Effects

There are some seasonal effects that are visible. The first is the weekly pattern which is described above. The second is that there has been a steady increase in volume over the course of 2.5 years, which is quite in line with the growth of the e-commerce market in general. Furthermore, there is a difference in volumes between different periods within the same year, as can be seen in Figure 6. This is due to the fact that at the end of the year there are many promotional actions and national holidays for which consumers make purchases.

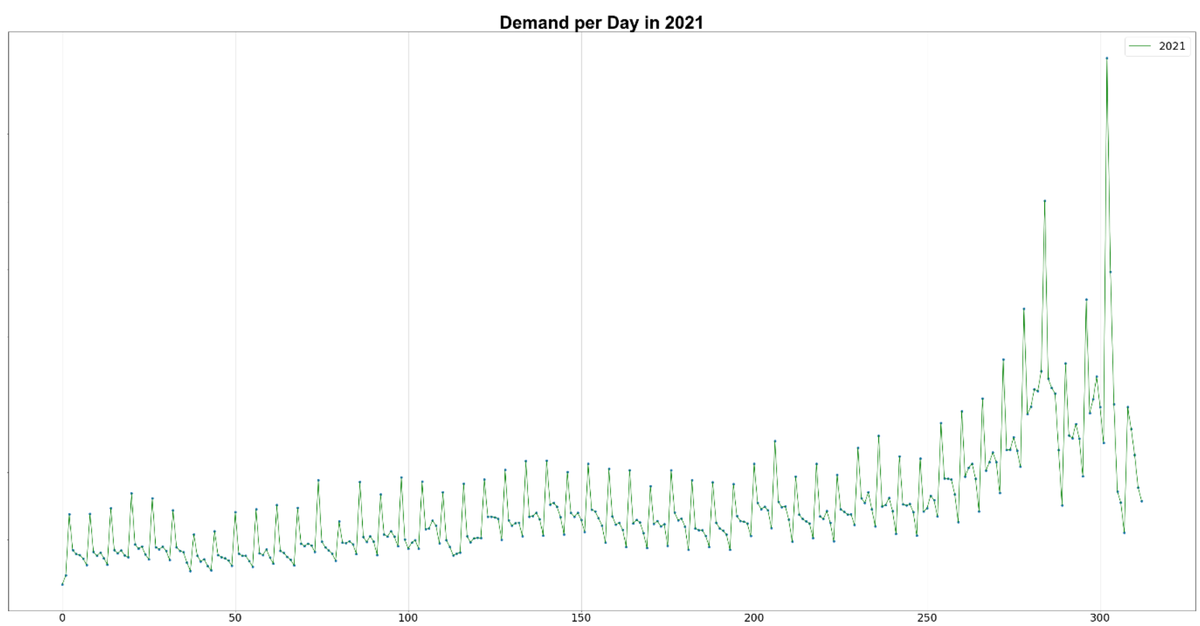

Figure 6 Difference in volume per day over the course of 2021[2]

2.2.3.4 Distribution of volume over the country

Lastly, we have investigated how the volumes are distributed over the country. We look at the level of 'postal code 4 areas. In the Netherlands, there are just over 4000 of these areas. Because the outlook is relatively similar we do not show the division for every weekday in every period for every logistical process. Therefore, we show the average volume of a Monday, where it is being collected, and where it needs to be delivered (left and right maps in Figure 7 respectively). For collection volumes, there is a relatively equal division of locations over the country with the exception of the northern areas. In the design of a network, the right map is more interesting. There are some areas, especially in the western and southern urban areas, where much more deliveries must be done in comparison to others. Choosing depot locations conveniently to those areas will likely save a lot of costs in the routes that need to be driven.

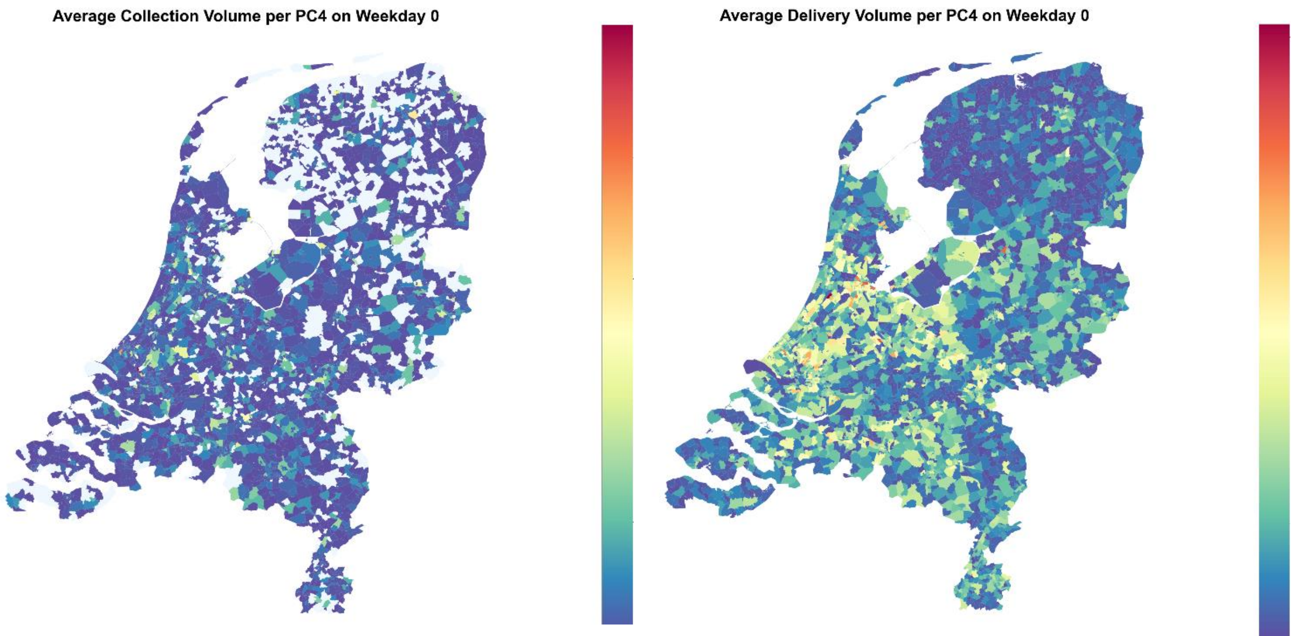

Figure 7 The distribution of average PPC volumes on Monday and the corresponding destinations (in parcels)[2]

2.3 LOGISTICAL PROCESSES

To better understand what type of logistical processes have to be performed by Company X, the multiple logistical processes need to be explained. As explained, the e-commerce party sells products that they have to collect from other retailers, these can be divided into two types: piecewise parcels (PP) and batch parcels (BP).

2.3.1 Piecewise Parcel Collection

Firstly, the group where items are collected on an item basis (PP). These retailers receive an order through the e-commerce company's web shop. The next day, in the morning, Company X will come to pick up all orders that have been placed the day before. This needs to be done in the morning since these orders need to be delivered the same day (the day after ordering). If the retailer is large, Company X will come to that retailer to collect the parcels, these are pick-up retailers. If the retailer is small, the agreement is that the retailer brings its parcels to a 'collection point' before a certain time, these retailers are indicated as "deposit retailers" since they "deposit" their goods at a collection point. These collection points are located within a driving time of X[3] minutes and can be, e.g., larger retailers that would have been serviced anyway. Once all parcels are collected, they are sorted by Company X and handed over to other parcel carriers who will do the delivery to consumer homes. This process is indicated as a piecewise parcel collection (PPC). An example of the network for the PPC process can be seen in Figure 8.

[3] *Due to confidentiality, we cannot disclose the exact numbers.*

2.3.2 Batch Parcel Collection

The second group are the retailers whose products are collected in larger batches and are then stocked in the warehouse of the e-commerce party; these are indicated as 'batch retailers'(BR). These retailers are responsible for making sure that their stock at the ECC warehouse does not run out. These retailers can put out a request to have Company X come and pick-up a replenishment. This involves larger boxes containing multiple items. The collection of this type of parcels is done after the collected parcels of the first type are sorted. This process is indicated as a batch parcel collection (BPC). An example of the network for the PPC and BPC process can be seen in Figure 8 and Figure 9.

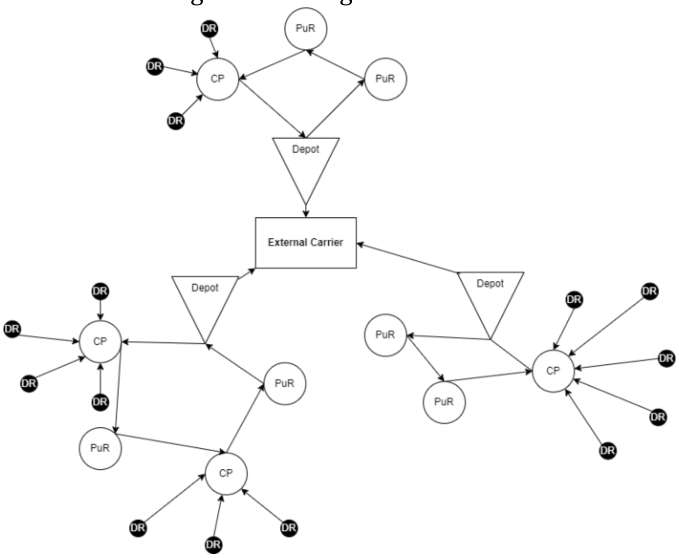

Figure 8 Schematic graph of PPC network

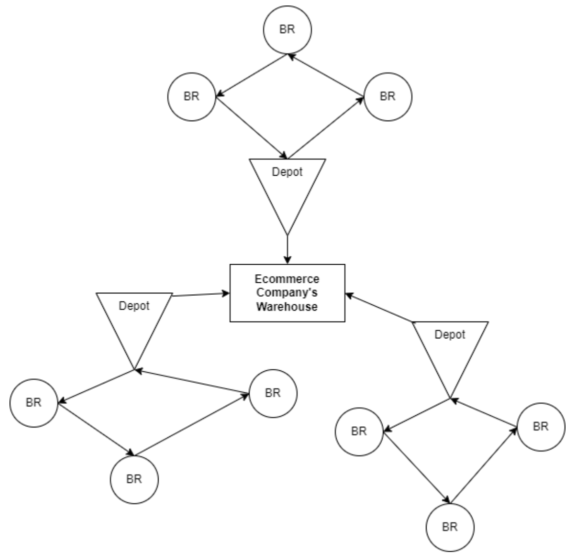

Figure 9 Schematic graph of BPC network

2.3.3 Sorting

When the PP collection is done, the parcels are sorted based on their destination. At the moment, the parcels are handed over to other parcel carriers, therefore the directions on which the parcels are sorted are the same as the depot locations of the third-party carrier. This sort is called the first sort (FS). The first sort needs to be done in a relatively narrow time window to be able to hand the parcels over to the third-party carrier so that the carrier has enough time to distribute them over the country and deliver them to customers.

2.3.4 Linehaul

When the parcels are sorted at a depot, they need to be transferred to other locations. For the PP collection, this is done by other parcel carriers who come to Company X's depots to collect the parcels there. For BP collection, the parcels must be consolidated at the depot from the different collection vehicles and transferred to the warehouse of the e-commerce company. Transport of parcels between depot locations is referred to as linehaul (LH).

2.3.5 Home Delivery

As indicated, Company X is looking into the possibility of also taking up the home delivery process (HD). This process includes the delivery of piecewise parcels to the homes of customers that have placed an order. These parcels can come either from retailers through the PPC process or from the e-commerce company's warehouse (ECW). Parcels are delivered to a depot, where they are sorted on route (second sort). From there they are delivered to customers' houses. Including home delivery will also require that Company X performs additional linehaul transportations during the day to move parcels to the right depots.

Table 1 An overview of the used definitions

Definition	Abbr.	Explanation
Piecewise Parcels	PP	Parcels that contain one single customer's order, destined for that customer's home.
Batch Parcels	BP	Large parcels, containing multiple similar items, that are destined for the e-commerce company's warehouse.
Pick-up Retailers	PuR	Large retailers (daily volume >X) that are serviced individually
Batch Retailers	BR	Large retailers that ship batch parcels are serviced individually. These retailers can also be PuR, depending on the portfolio of products they offer.
Deposit Retailers	DR	Small retailers (daily volume <= X) that bring (deposit) their orders to a collection point
Collection Point	CP	A PuR or DR where parcels of DRs are aggregated to achieve full(er) truckloads
E-commerce Company's Warehouse	ECW	Warehouse from where the orders containing EC's own products or products that have been sent by Batch Retailers.
Depot	-	A location where collection or delivery routes start and end as well as where parcels are sorted and consolidated.

Table 2 An overview of the different logistical processes

Definition	Abbr.	Explanation	Parcel Type	Time Window[4]
Piecewise Parcel Collection	PPC	Collection of PP Parcels	PP	Morning
Batch Parcel Collection	BPC	Large parcels containing multiple similar items, destined for the e-commerce company's warehouse.	BP	Afternoon
Linehaul	LH	Transportation of parcels in-between depots and/or ECW.	PP/BP	Afternoon & Evening[5]
Home Delivery	HD		PP	(Early) Evening
First Sort	FS	The process of sorting collected parcels based on their destination's depot.	PP	Early Afternoon
Second Sort	SS	The process of sorting parcels based on delivery routes	PP	Late Afternoon

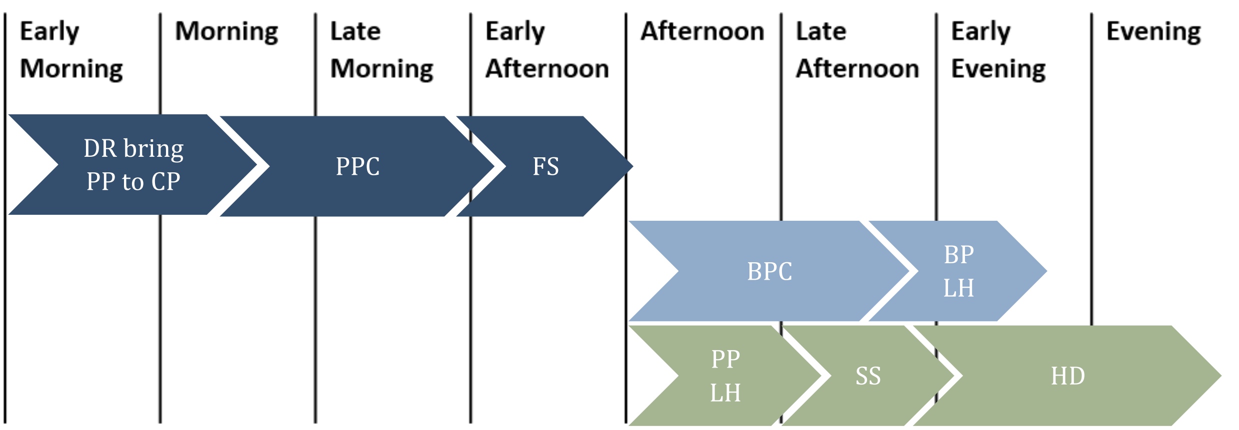

Figure 10 Overview of different processes over the day (for abbreviations, see Table 1 and Table 2)

[4] *Due to confidentiality reasons, the exact time windows are not mentioned.*
[5] *Evening for BP and potentially in the afternoon in case Company X also will perform Home Delivery.*

2.4 NETWORK DESIGN FOR COMPANY X

As explained earlier, Company X already has had an analysis done to determine how they should roll out their logistical network over the Netherlands. In this section, we provide a short outline of how that network design has been done. Furthermore, we address the KPIs that Company X has for determining the quality of a network design.

2.4.1 Network elements

Before we explain the process of how the current network design has been made, we must first touch upon the decisions that Company X has to make when designing the network.

- The first is the decision of which retailers should become collection points. As explained in Section 2.3.1, the smallest retailers should deposit their goods at a collection point. To have the most efficient routes, this collection point must be chosen as conveniently as possible.
- The second element is choosing the depot locations, these speak for themselves and ideally have to be chosen such that the daily operational cost is as low as possible today as well as in a few years.
- When all locations in the network have been set, the third element to decide upon is the routing. Routing can, depending on what Company X will do in the future, be split into collection routes and delivery routes.
- The last element is the linehaul. In case Company X is also going to perform delivery routes, the collected parcels need to be sorted at the depot and transported to other depots from where they will be delivered. This inter-depot shipment of parcels is the linehaul. Note that the linehaul volumes are the result of the chosen service areas, as the route of a parcel is bound by the depot allocations of its origin and destination postal zones.

2.4.2 Current network design

The current plan for the logistical network has been determined using some input parameters and decision rules. As mentioned, Company X is a start-up, therefore this modelling has been done based on historical data from ECC and assumptions. In this section, we provide a high-level overview of the most relevant steps.

The choice for collection points is done by means of a clustering rule based on the weights of shipments and the requirement that any retailer should be within a driving range of at most X[6] minutes from a collection point. These clusters are then split into subclusters based on the maximum volume or number of retailers allowed at one collection point. For each subcluster, the closest pick-up retailer is appointed as a collection point.

Secondly, the depot locations are chosen based on spatial distribution and the weight of the delivery volume of PC4 areas. This means that implicitly there is a direct link between each depot and every PC4. For depot sizes, a predetermined set of sizes is used from which can be chosen when locating a depot.

Thirdly, the routing. For the piecewise parcel collection (PPC), two scenarios of a day are simulated, one 'regular' day and one for the peak season. For the batch parcel collection (BPC), one average day scenario is used. The routing is composed of an approximation of travel time from a depot to a PC4 area, the travel time and stop time within a PC4 area, and the time to travel from one PC4 area to another. For the travel time within a PC4 area, an approximation is used based on the density of stops and the level of urbanization of that area (the higher the urbanization, the lower the average speed a vehicle can drive). Note that this routing is not used to improve the network design, but merely to check whether the network resulting from steps one and two would be cost-effective.

Lastly, the linehaul for PPC and BPC. The network is assumed to be fully interconnected, meaning there is a direct linehaul connection between every depot in the network. Furthermore, it is assumed that the linehaul volumes from depot A to depot B are proportional to the delivery volume weights of the corresponding areas.

[6] *Due to confidentiality, we cannot disclose the exact numbers.*

All in all, the initial design is a good starting point for Company X. Because of the demand-weighted approach, the clusters and depot locations will probably be placed in the proximity of the best locations as these will be the locations with the most routes. However, it also becomes clear that the design is based on many assumptions and approximations, especially in the routing part. Therefore, this initial design can be improved further in order for Company X to have the most efficient network design.

2.4.3 KPIs for decision-making used in the current network design

Then, it is important to know which key performance indicators (KPIs) have been used in the network design. Although there are multiple, the KPIs can be divided into three main categories: costs, capacity, and productivity. Costs include fixed costs for vehicles, depot locations, and overhead as well as variable costs like wage and driving costs per km. Capacity indicates the volumes that on average end up in a vehicle. Lastly productivity, this relates to the time spent per stop and route, i.e., the time required to perform the processes.

2.5 TRANSPORT RESOURCES

In the previous sections we have addressed almost all relevant information for designing the network of Company X. However, we still miss two inputs that we need to clarify. These are the depot types, sizes and costs, and the vehicle types, capacity, and operating costs.

2.5.1 Depots

Depots can come in various sizes, types, and locations that all influence the costs of the depot. Furthermore, the depots of Company X do not exist yet, which makes it even harder to make a good estimation of size, type, location, and thus costs. Therefore, we have made some approximations to determine which locations to take as input.

First, the depot costs, multiple components can be distinguished: rental costs, one-time investments (capex) like a sorter, operational costs, and personnel costs. Based on internal documents of Company X, when all these costs components are added up to account for a cost per year for different types of depots. Because at this point it is not possible to know how many depots of each type are required, an estimation is done on the number of depots of each type based on the spread of historical sales volumes over the country. Adding up these total yearly costs and dividing those by the number of working days in a year we get an average daily cost of opening a hub of €600.

Secondly, the locations and number of hubs that will be considered in the model need to be determined. To do this, we took a publicly available dataset of all commercial areas (Dutch: industrieterrein) in the Netherlands, which was published by the organisation of Dutch provinces (Interprovinciaal Overleg, 2022). From this dataset, we selected all areas that were not yet full and had a good enough road connectivity (up to 15 minutes from the highway). This resulted in a subset of around 950 possible locations out of 3800 commercial areas. Because this dataset does not include postal zones but only the city name and the name of the commercial area, we used the geopy package for python to estimate the postal zone of each commercial area. Consecutively, we determined a list of hub locations that ensured nationwide coverage.

We created this list by calculating for each commercial area how many postal zones were within a 60-minute drive time, for which we used a dataset of driving times between each postal zone in the Netherlands that was created for the VPRO project 'Nederland van Boven' which is based on the Dutch 'Nationaal Wegen Bestand' (GeoDMS, 2019). The first entry on the list was the area that had the largest coverage. Next, we calculated for each commercial area how many postal zones they would be able to reach within 60 minutes and chose to add the area that could add the most postal zones. This process is repeated until no more postal zones could be added, resulting in 15 hub locations.

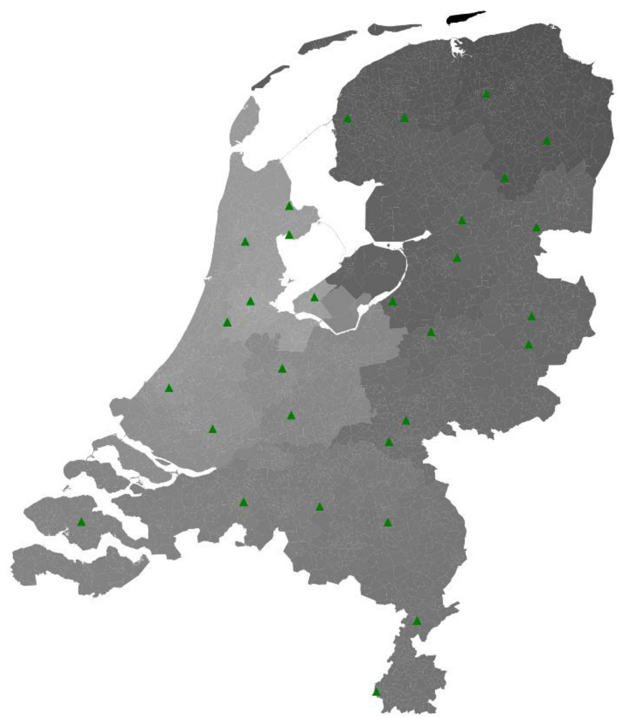

Figure 11 Candidate hub locations

Since this number is not very large and does not provide much flexibility in choosing hub locations, the list is expanded. For each hub location in the list, the three closest other hub locations are chosen. For these three hub locations, we chose the location that was closest to the centre point between the three hub locations. This resulted in 15 additional hub locations that provided a more even spread of candidate hub locations over the country.

Lastly, some candidate hub locations were positioned very close to each other. Those were moved manually to increase the spread of locations even more. The final selection of 30 hub locations is shown in Figure 11.

2.5.2 Vehicles

The characteristics of the vehicles, like capacity and driving range, influence the number of stops in the routes and thus the distances and durations of the routes. Furthermore, there is a fixed cost per usage of a vehicle and a cost per driven kilometre. Lastly, the driver of the vehicle needs to be paid a salary per hour. For this research, two types of vehicles are considered. The first is the large van (LV), which is used for the collection of piecewise and batch parcels. The second is the small van, which is used for home delivery because it is easier to use in highly urban areas.

Table 3 Vehicle characteristics

Shielded for the public version

** Since batch parcels are bigger, less fit in the van*

In reality, for each process, a mix of different vehicle types could be used to get the lowest possible routing costs. However, for this research, it is assumed that for each process only one vehicle type can be used.

2.6 CONCLUSIONS ON THE CONTEXT ANALYSIS

In this chapter we have answered the research question '*How is the current logistical network design done for Company X?*'. Firstly, we established an overview of the (expected) workload for Company X by looking at trends in the Dutch e-commerce market in the Netherlands. Important observations are the high yearly growth rates, which imply further growth for the coming years, and the high demands Dutch consumers have on the delivery process of their orders. Furthermore, we have seen that the number of retailers selling through the ECC has been increasing rapidly over the last few years. These retailers are spread over the Netherlands, with many postal zones that contain retailers. As can be expected, the number of retailers in urban areas is higher. Although the number of retailers is very high, the number of pickup retailers (200) is relatively low at only 6% of the total. Also, there are some clear patterns visible in daily workload within weeks and years. Within a week, Monday is the day with the highest workload, which is expected since orders placed during the weekend are also included. Furthermore, a great peak can be seen at the end of the year due to several holidays.

Secondly, we looked at the different logistical processes that Company X does and possibly will perform. There are five processes that are performed during the day, that together form the logistical fulfilment of piecewise and batch parcels. Especially, the processes related to the piecewise parcels have a great impact because the parcels are collected and delivered on the same day.

Thirdly, we looked at how the current network design has been performed. The choice for depot locations and service areas has been made on clustering rules and spatial distribution of postal zones to the depots weighted by the respective volume of each postal zone. The driving distances are approximated by the number of stops and level of urbanization of an area. For the linehaul, volumes are assumed to be proportional to the delivery volume of each hub. In this design, three KPI categories are used: costs, capacity, and productivity. This approach is a good starting point because of the demand-weighted clustering. However, this approach is based on many assumptions and approximations. Therefore, a more comprehensive approach will enable Company X to build a more efficient network.

Lastly, we constructed a set of 30 possible hub locations to consider in the model. These are based on all industrial/commercial areas in the Netherlands. Furthermore, two vehicle types are distinguished that are used in the collection and delivery processes.

3 LITERATURE REVIEW

In this chapter, we research the literature to answer the second research question '*Which models exist in the literature that optimize transportation networks with depots and hubs?*'. Answering this question will enable us to find a way to model the logistical network design of Company X. First, we will look into the literature on network design, after which we will address more specific parts of network design literature that are relevant to our problem. Lastly, we address methodologies to solve real case problems that are described in the literature.

In the case of Company X, the design of a logistical network consists of four elements. Firstly, you need to determine the number, sizes, and locations of depots. Secondly, you need to assign locations/areas to depots to create service areas. Thirdly, you need to determine the line-haul structure. Lastly, you need to determine the optimal routing from all the depots to all locations that need to be served (within this research this is split into collection and delivery routes).

In the literature, we found multiple optimization problems that relate to one or more of the above-described network elements. The first three elements together make up the Hub Location Problem, which is described in Section 3.3. The routing optimization is described by the well-known Vehicle Routing Problem, which is left out of scope for this literature review. Combine those two problems and you have the Hub Location Routing Problem, which is described in Section 3.4. A graphical representation of how the network elements relate to the different optimization problems is shown in Figure 12.

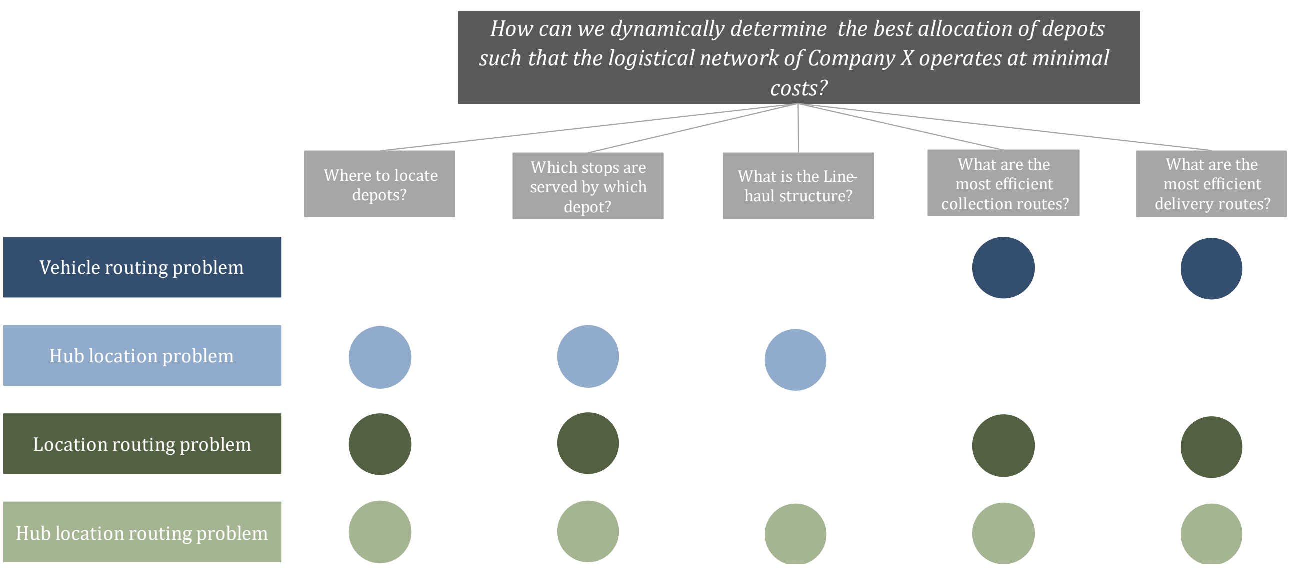

Figure 12 An overview of how different network design elements are covered in theoretical known problems.

3.1 NETWORK DESIGN PROBLEMS

Network design problems arise whenever optimal choices have to be made that can be represented conceptually as the selection of a subset of links in a graph. Typically, these optimal choices are the result of complex trade-offs between various types of costs and constraints (Crainic et al., 2021). In particular, most network design problems involve fixed costs associated with link selection and variable costs associated with flows (of people, goods, information, etc.). Because of their combinatorial nature and the complexity of their objective functions and constraints, network design problems are inherently difficult (most of them are NP-hard) (Crainic et al., 2021).

There are a few important characteristics that determine the type of network design problem. The first fundamental distinction in network design problems is whether the demand can be represented as one commodity (possibly with multiple origins and multiple destinations) or as multiple commodities that need to be differentiated. The second characteristic is whether capacities are taken into account. And the third characteristic is the cost structure, the presence of fixed design and variable transportation costs introduce complex trade-offs. Making finding optimal solutions more complicated (Crainic et al., 2021).

3.1.1 Fixed Charge Network Design Problem

At the basis of network design problems lies the fixed charge network design problem (FCNDP). The most basic case is where one single commodity is transported from an origin to a destination through a graph over arcs and nodes. The multicommodity variant has multiple commodities and thus multiple origin and destination nodes (Crainic et al., 2021). Based on this model, further expansions can be made which we will not address due to lack of relevance to our case.

The single commodity formulation for the FCNDP, given by Crainic et al. (2021), is as follows. Let $G = (N, A)$ be a directed graph, where N is the set of nodes and $A \subseteq N^2$ the set of potential arcs. A limited flow capacity $u_{ij} > 0$ is associated with each arc (i, j). The network design problem consists of selecting a subset of arcs from A to satisfy a given demand at the minimum total cost. The demand to satisfy is defined at the nodes of the graph, which are partitioned into three subsets: N^o, the set of origin (source) nodes, N^d , the set of destination (sink) nodes, and N^t , the set of transhipment (intermediate) nodes. Each origin $i \in N^o$ has a supply (availability) $w_i > 0$ of the given commodity, each destination $i \in N^d$ has a demand (request) $w_i < 0$ of the same commodity, while each transhipment node $i \in N^t$ has neither availability nor request, i.e., $w_i = 0$. The net supply across any set $S \subseteq N$ is defined as $W(S) \equiv \sum_{i \in S} w_i$. We assume that demand is balanced, i.e., $W(N) = 0$.

3.1.2 Service Network Design Problem

In their book, Crainic et al. (2021) introduce the Service Network Design problem as a starting point for network design problems with applications in transportation and logistics. In this problem, the supply side of a transport system is designed in order to meet demand in a high-quality and profitable way. Service is understood to be a vehicle, convoy, or train which moves through the network transporting people or freight loads. Service Networks are especially applicable in the context of consolidation-based transport, often organised in hub-and-spoke networks (Section 3.2). One can think of postal and small-package transportation companies, less-than-truckload (LTL) motor carriers, railroads, ocean/maritime liner navigation companies, and land- and water (coastal, river, etc)-based intermodal carriers. (Crainic & Hewitt, 2021).

3.1.3 Hub Network Design Problems

One subset of these applications of service networks is when the network design involves the location of consolidating hubs, hub network design. A distinguishing feature of hub networks is the use of transhipment, consolidation, or sorting points for commodities, called hub facilities, to connect a large number of origin/destination (O/D) pairs by using a small number of links (Contreras, 2021). Hub networks can be seen as hierarchical networks which, in their most basic form, contain two levels: an access-level network connecting O/D nodes to hubs, and a hub-level network connecting hub nodes between them (Contreras, 2021). These hub network design problems were first studied from a facility location perspective, but now there is a specific subset that is concerned with hub locations, and their interactions, as the main decision. These are the hub location problems (Section 3.3). Recently, literature is published on combining hub location problems with collection, transfer, and delivery routing (Contreras, 2021).

Hubs are facilities of various types where one or more transportation modes interact. They are intermediate facilities that perform switching, sorting, connecting, and consolidation/break-bulk functions for traffic (e.g., passengers, freight, or information) between many origins and destinations (Alumur et al., 2021; Contreras, 2021). We can think of hubs in many-to-one and one-to-many systems as consolidation points that allow linehaul and local operations to be decoupled (Daganzo, 2005).

3.1.4 Hierarchy in network decisions

In transportation planning, there is a hierarchy in decision making which includes strategic, tactical, and operational decisions (Magnanti & Wong, 1984). Strategic decisions are long-term decisions related to the infrastructure of transportation networks, e.g., the building of highway infrastructure. Tactical decisions are those concerned with the effective utilization of infrastructures and resources of existing transportation networks (in contrast to the acquisition of these resources). Lastly, operational decisions are short-term decisions, mostly related to traffic flow control, demand management or scheduling. Given our research question, we are dealing with decisions on a strategic, locating depots in a network, as well as tactical level, determining the service area for each installed depot.

3.2 HUB-AND-SPOKE NETWORKS

Hub-and-spoke (HS) networks are used in various network systems, for example, airline systems, and telecommunication systems. In the context of postal service, this network structure is commonly used (Wu et al., 2022). An HS network is characterized as a network type, used in various many-to-many systems, with hubs and non-hubs, where large volumes are transported between hubs to achieve economies of scale and non-hub nodes, are used to distribute volumes, passengers, or information over the different regions (de Camargo, 2013; Wu et al., 2020). It is known that a hub-and-spoke structure is suitable for those delivery systems where it is expensive or impractical to dedicate exclusive transport links to each origin-destination pair and it is a well-known configuration in postal systems (Çetiner et al., 2010).

Hub-and-spoke networks can offer some benefits in multiple cases in which we can have different possibilities for the number of origins, destinations, and commodities that have to be transported through the network (Lopes et al., 2016; Pandiri & Singh, 2021).

- One-to-Many: one origin has to serve different destinations.
- Many-to-One: many origins have to serve one destination.
- Many-to-Many: many origins have to serve many destinations.

As an illustration for the one-to-many case, it is assumed that each destination demands a specific number of items from each one of the origins and that these cannot be substituted for one another. That is, we are dealing here with what normally is referred to in the network optimization literature as a multi-commodity problem (Daganzo, 2005).

In such HS networks, the consolidation of flows increases the traffic density in some (or most) route segments. For a transportation setting, this greater traffic density allows using larger and more cost-efficient vehicles (e.g., aircraft or trucks) – with appropriate trip frequencies. The reduction of unit costs comes from sharing fixed costs over more units of demand (e.g., passengers), and possibly from using vehicles with lower variable costs. Additional benefits of hub networks can come from increasing the frequencies of service on links (as a result of higher traffic density), and a better traffic balance across the network. Further, hubs can concentrate administrative and technical resources, reducing the investment, operational costs, and inventories. (Alumur et al., 2021)

3.2.1 Single and multiple allocation

Originally, HS networks were assumed to have a direct connection between every hub pair; no two non-hub nodes could be serviced by a direct link; non-hub nodes could be directly connected to only one of the installed hubs, i.e., were single allocated, or could send and receive flow through more than one of the installed hubs, i.e., were multiple allocated; an origin–destination demand was routed through one or at most two hubs, and economies of scale on inter-hub connections were represented by a discount factor (de Camargo et al., 2013).

However, over the years, these assumptions have been adjusted and adapted to cope with different problems and applications such as the design of a ring star hub network; the location of inter-hub connections instead of just hubs; the minimization of the maximum travel time between any origin–destination pair; allowing at most three hubs on any given route between pairs of origin–destination or incomplete inter-hub connections; but also minimizing the maximum travel time between any origin–destination pair; projecting star-star like hub networks; having a tree structure connecting all installed hubs; and partitioning the points into sub-networks, locating at least one hub in each sub-network and routing the flow at minimum cost. (de Camargo et al., 2013)

3.3 HUB LOCATION PROBLEM

The hub location problem (HLP) was first introduced by O'Kelly in 1987 (O'Kelly, 1987). Broadly speaking, HLPs consist of locating hub facilities, designing the hub network and determining the routing of flows (individuals, goods, and information) from an origin through the network to their destination, while optimizing a cost-based or service-based objective (Alumur et al., 2021; Campbell, 1994). Hub networks are often used in transportation, telecommunication, and computer systems to efficiently route commodities between many origins and destinations (Alumur et al., 2021; Contreras, 2021). The HLP is part of the class of NP-Hard problems (Contreras, 2021).

The primary advantages of hub networks stem from 1) lower movement (i.e. transportation or transmission) costs from consolidated flows that exploit economies of scale, especially between hubs; 2) reduced costs from establishing a sparser network to connect many dispersed origin–destination (O–D) pairs; and 3) better service from allowing more frequent connections due to the consolidated flows (Alumur et al., 2021).

Network hub location models include origin/destination (O/D) nodes, hub nodes (to be located), an access network to connect the non-hub O/Ds to hubs, and an inter-hub network connecting the hubs. On top of that, a key modelling decision is whether to assume a topology of the access-level, connecting O/D nodes to hubs and hub-level networks, connecting the hub nodes (Alumur et al., 2021; Contreras, 2021). Demand in HLPs is a request for service for a specified O-D pair where something (e.g., information or freight) is transported from an origin to a destination, possibly via one or more hubs. These hubs function as intermediate points delivering some desired service (e.g., sorting) (Alumur et al., 2021).

Hub location-allocation has various application areas in digital data service, telecommunication networks, air transport, freight transportation, postal services, public and urban transportation, rail transport, and emergency services. (Alumur et al., 2021; Basirati et al., 2020; Contreras, 2021)

3.3.1 Main model elements

An HLP is generally described as a graph $G = (N, E)$, where N is the set of nodes representing the origins and destinations (non-hubs) of flows as well as the set of potential hub locations, and E is the set of edges. For each node pair (i, j) there is an amount of flow to be routed (W_{ij}) and a distance (d_{ij}), both non-negative, from origin i to destination j. Fixed costs are incorporated with each node for opening a hub, and with every edge for activating inter-hub arcs. These hub arcs connect hub nodes and have a unit flow cost of $\alpha \cdot d_{ij}$, where α is a discount factor to reflect economies of scale. Furthermore, each O-D path has a collection leg from the origin to the first hub and a distribution leg from the last hub to the destination (Contreras, 2021).

Most proposed models for the HLP are mixed-integer linear programming (MILP) models, yet in recent literature also non-linear formulations have been proposed more often (Basallo-Triana et al., 2021). Within the formulations, different types of transportation variables can be defined. Flow-based variables use three index continuous variables to represent the flow from an origin routed through two hubs. Path-based variables use four subscript binary variables for selecting the route for each O-D pair through the network (Basallo-Triana et al., 2021). Flow-based variables have been shown to be superior in terms of solution time according to Alumur and Kara, yet a benders decomposition method turned out to be faster for the hub location routing problem formulation (Section 3.4) (Basallo-Triana et al., 2021).

3.3.2 Assumptions

Classic HLPs rely on some main assumptions. First, demand flows have to be routed through one or at most two hubs, implying that a non-hub node is directly connected to at least one hub facility. In other words: unit shipment from a location to a customer is independent of the route taken to visit the customer. Second, the hub network is considered to be fully interconnected. Third, a constant discount factor representing economies of scale is applied to the unit transportation cost of inter-hub connections. In several applications, these assumptions are reasonable and provide a good approximation of reality; in others, they may lead to suboptimal solutions. Furthermore, delivery costs depend only on the sum of the product of the unit shipment cost from the location to the customer and the number of units delivered to that customer (Alumur et al., 2021; O'kelly, 1987; Real et al., 2021)

Demand in HLPs is usually implied to have a common time period for all O-D pairs and HLPs are generally solved for a single time period which is presumed to repeat. However, in reality, demand is not static but is a rate of some unit per time and is generally dynamic. Furthermore, operations in a hub network take relatively little time whereas travelling takes much more time. In the literature on hub locations, little attention has been paid to the time dimension (Alumur et al., 2021).

Hub location problems were developed especially for the linehaul design problem. Classical location-allocation problems, like the HLP, also ignore tours when locating facilities. Most HLPs use a set of assumptions that simplify the design and routing decisions to the point of being completely determined by the allocation decisions of O/D nodes to hubs. This may lead to incorrect pickup and delivery costs (Contreras, 2021; Wasner & Zäpfel, 2004).

3.3.3 Topologies & Allocation Strategies

In the design of a hub network, there are two levels, the access level and the hub level. On both levels, choices can be made on how the nodes are (inter)connected which will influence the way the final network will look. On the access level, this concerns allocation strategies. On the hub level, there are multiple network topologies that could be applied.

3.3.3.1 Allocation Strategies

In HLPs, three allocation strategies for non-hub nodes to hubs have been investigated: 1) single allocation, each node is assigned to one hub; 2) multiple allocation, nodes can be assigned to more than one hub; and 3) r-allocation, each node can be assigned to at most r-hubs. Multiple allocation has been the preferred strategy, although, in hub location routing context, authors chose single allocation more often (Alumur et al., 2021; Basallo-Triana et al., 2021; Contreras, 2021).

3.3.3.2 Topologies

Within HLPs, the topology is of great importance. Next to fully interconnected where each hub is directly connected to all other hubs, we describe four different topologies: star-star, cycle-star, tree-star, and line hub networks, see Figure 13. Star-star hub networks consist of a set of hub nodes connected directly to a central hub node and every non-hub node is connected to a hub node, creating a star network on access and hub level. Applications can be found in satellite communication or cargo delivery networks. Tree-star networks consist of a set of hub nodes connected via a spanning tree. Each non-hub node is assigned to exactly one hub. Applications can be found in digital data services or train networks. Cycle-star networks consist of a set of hub-nodes connected by a set of hub arcs that form a cycle, each non-hub node is connected directly to one hub. Applications can be found in telecommunication networks or rapid or public transit systems. Lastly, hub line networks consist of a set of hub nodes that are connected through a path (line) and non-hub nodes can be assigned to one or more hubs. Applications arise in public transportation or road networks (Contreras, 2021).

3.3.4 Improvements

Alumur et al. (2021) define nine themes on which the HLP models can be improved compared to existing literature: 1) Better model economies of scale; 2) Incorporate time in HLP models; 3) Consider more sophisticated objectives and multiple criteria; 4) Relate to real-world problems; 5) Integrate hub location with other problems; 6) Incorporate nature of the demand; 7) Use real data; 8) Obtain insights from the results; 9) Use the best solution approaches – and develop new ones.

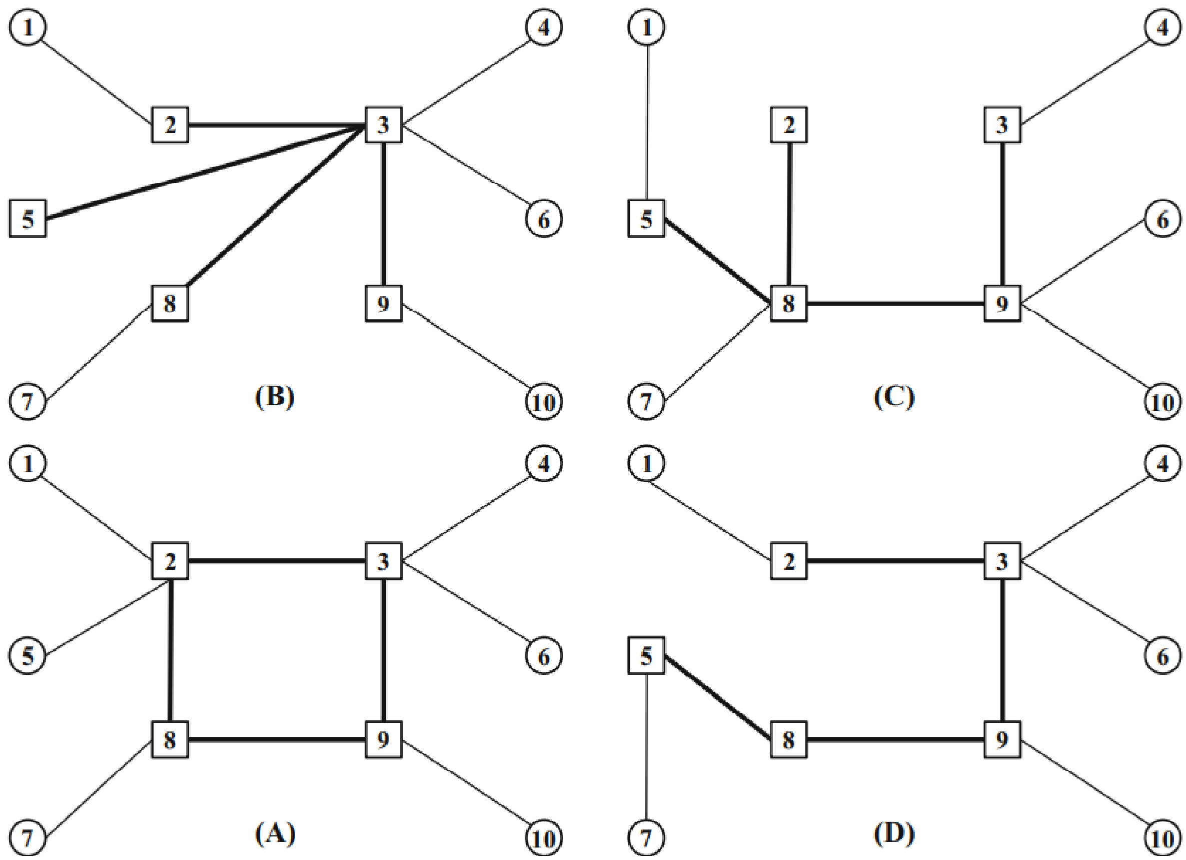

Figure 13 Structure of (A) Cycle-star, (B) Star-star, (C) tree-star, and (D) line hub networks (Contreras, 2021)

3.4 HUB LOCATION ROUTING PROBLEM

The hub location routing problem (HLRP) is a problem consisting of all decisions related to the locating of hubs, generating multiple-stop local routes for the non-hub nodes, allocating non-hubs to the installed hubs, and routing the flow of many origins to many destinations at a minimal cost, by which it aims to meet the demands by moving the individuals, goods, and information. The HLRP essentially combines the hub location problem with the vehicle routing problem (Basirati et al., 2020; Campbell, 1994; Çetiner et al., 2010; de Camargo et al., 2013; Lopes et al., 2016; Nagy & Salhi, 1998; Wasner & Zäpfel, 2004). The HLRP can be thought of as a strategic or tactical problem (Çetiner et al., 2010).

Since hub location and routing influence each other, it is imported to consider the two problems simultaneously (Basirati et al., 2020). Hub location-routing models include constraints on routing aspects to model multi-stop collection and distribution routes at hubs. Extending HLPs to hub location-routing problems allows multi-stop routes in the access network connecting hubs and non-hubs (Alumur et al., 2021).

Just as with the HLP, within the class of HLRPs, there is a distinction between single-allocation and multi-allocation. Single-allocation HLRPs allow non-hub nodes to be allocated to only one hub node, whereas multi-allocation HLRPs allow for non-hub nodes to be served by multiple hub nodes (Wu et al., 2022). On top of that, there are different types of networks. One destination receives flow from different origins, the many-to-one case; or one origin node can supply many destination nodes, the many-to-one case; or thirdly many origins can supply many destinations, the many-to-many case (Lopes et al., 2016; Pandiri & Singh, 2021).

Hub nodes are special facilities that serve as connections between origins and destinations, these are also called non-hub nodes. They act as consolidation centre that bundles the quantities of parcels of certain demand points to achieve economies of scale for less-than-truckload transports. Pickup and delivery routes start and end here (de Camargo et al., 2013; Lopes et al., 2016; Ratli et al., 2022; Wasner & Zäpfel, 2004).

In general, three assumptions are considered in classic hub location-routing problems including (1) the network structure is based on a complete graph, (2) interactions among hubs include the discount factor α, (3) direct routes among non-hub nodes are forbidden (Basirati et al., 2020; Karimi & Setak, 2018).

3.4.1 Location Routing Problems

In the previous sections, we have seen two problems that describe two sets of decisions in the challenge of designing a logistical network. There is a third, well-researched, type of mathematical problem that touches upon network design, namely the set of location routing problems (LRP) which were introduced by Watson-Gandy and Dohrn (1973). Location routing problems are characterized by the search for the optimal number and locations of distribution centres, simultaneously with the vehicle schedules and distribution routes so as to minimize the total system costs (Min et al., 1998). Generally, in LRP models the facilities are not connected and clients do not exchange flow (de Camargo et al., 2013).

For readers interested in LRPs, we refer to Drexl & Schneider (2015), Zhou et al. (2016), Zhou et al. (2018), and Mara et al. (2021) for reviews of the LRP literature.

3.4.2 Literature on the Hub Location Routing Problem

The (many-to-many) hub location routing problem (HLRP) is an expansion on the classical LRP. The hub location routing problem was introduced first by Nagy and Salhi (1998) as the many-to-many location routing problem (MMLRP). They proposed a model and hierarchical heuristic solution approach in which they allowed pick-up and delivery on different occasions, while also applying capacity and maximum distance constraints (de Camargo et al., 2013; Nagy & Salhi, 1998).

In this section, we will explain the literature on the HLRP structured by solution method. A complete overview of the different types of solution methods can be found in Figure 14.

3.4.2.1 Exact Approaches

Due to the nature of the problem, exact approaches have been applied to small datasets of up to 100 nodes. The first exact solution approach was proposed by Catanzaro et al. (2011), who proposed the partitioning HLRP (PHLRP). This formulation consists of partitioning a network into sub-networks that each contain at least one hub. To solve the PHLRP, a MIP model was proposed. To decrease the required time to solve the MIP, a branch-and-cut algorithm was proposed. This showed to successfully reduce the required time to solve instances. However, only instances of up to 20 nodes have been tested.

Furthermore, de Camargo et al. (2013) introduced a new formulation to the many-to-many hub location routing problem with a total cost minimization function. This study has focused on the same shipping and discount coefficient and simultaneous pickup and delivery. However, there are no capacity constraints, and each non-hub node is met once. They proposed a benders decomposition approach to determine exact solutions.

Lastly, Rodríguez-Martin et al. (2014) also proposed a branch-and-cut algorithm for the HLRP with a fixed number of hubs, one vehicle per hub and a constrained route length. By using families of valid inequalities and separation algorithms, subtour elimination, and capacity constraints their branch-and-cut algorithm was able to solve problems with sizes of up to 50 nodes (Lopes et al., 2016).

3.4.2.2 Heuristic Based Approaches

After Nagy and Salhi introduced the HLRP and proposed a heuristic, the next authors to propose heuristic approaches for the HLRP were Wasner & Zäpfel (2004) formulated a multi-depot hub-location vehicle routing model to determine the optimal location of depots and one central hub for a mid-sized postal carrier in Austria. They develop a seven-step local search algorithm with four feedback loops that was able to reduce the network cost by 14%. In their model they took all Austrian postal zones (2042) into account, making it relevant for this research.

Çetiner, Sepil, and Süral (2010) built further upon that work and formulated an HLRP model to determine the locations of sorting locations required in a network that services 79 cities for the Turkish Postal Services. To solve their HLRP formulation, they created an iterative two-stage heuristic. In the first stage hub locations are determined and postal offices are allocated to hubs. In the second stage, the routes within hub regions are determined.

Seto et al. (2021) considered a formulation of the HLRP that includes time windows for pickup and delivery, which makes it interesting for this research and allows for direct delivery between regions next to the 'regular' linehaul. Furthermore, they created a three-step network design algorithm that uses hierarchical clustering to deal with the uneven distribution of commodities over rural and urban areas. The first step of the heuristic is to determine the commodities that are delivered via direct delivery. The second step is to determine the regions and the hub location within each region. The third step is to determine the delivery routes and schedule (which are determined using an insertion heuristic). They were able to calculate solutions to problems of up to 50 nodes within 3 minutes.

3.4.2.3 Population Search-Based Approaches

Sun (2013) considered a pre-determined number of hubs and proposed a pHLRP with simultaneous pick-up & delivery and constraints on the depot and vehicle capacity. An ant colony optimization (ACO) was proposed to solve instances of 100 and 200 non-hub nodes.

Mokhatari & Abassi (2014) suggested a combination of variable neighbourhood search (VNS) and particle swarm optimization (PSO) meta-heuristics to construct a multiple hub transport network. It was assumed that the hubs and vehicles have an infinite capacity that is far beyond the existing, realistic situation in hub location-routing problems. The result of computational experiments demonstrated the superior performance of the proposed solution method as compared to existing algorithms.

Rieck et al. (2014) presented a variant of the many-to-many location routing problem from the timber-trade industry in which three layers (origin, hub, destination) are taken into account. For large-scale instances, they created a fix-and-optimize heuristic and a genetic algorithm (GA).

Bostel et al. (2015) applied the HLRP to a postal service system with simultaneous pick-up and delivery. They proposed a MILP and memetic algorithm (MA), a meta-heuristic combining a genetic algorithm with a local search to reduce premature convergence, which was able to solve instances of up to 100 nodes. They identify their formulation as the capacitated single allocation hub location routing problem (CSAHLRP). Lopes et al. (2016) proposed a biased random key genetic algorithm (BRKGA) for their p-location hamiltonian cycle problem, yet it performed worse than their variable neighbourhood descent.

Ghaffarinasab et al. (2018) proposed a continuous approximation approach to the planar HLRP in which demand is uniformly distributed over a polygon service area, so there are no nodes to visit. This approach determines new solutions using an equation and keeps iterating until the hubs in a solution have stayed within a distance v apart from each other for a set number of iterations. Compared to existing algorithms for the problem, their two proposed solution approaches, the iterative weiszfeld-type Algorithm (IWA) and particle swarm optimization (PSO) perform better by providing good quality solutions quickly, i.e., 650 seconds for 10 polygon areas for the PSO.

Yang et al. (2019) provide a MILP formulation and memetic algorithm (MA) for the single allocation HLRP with independent pick-up and delivery and capacitated hubs. Their MA performed well on the benchmark instances (up to 100 nodes and 10 hub locations) compared to the CPLEX and in a reasonable time.

Basirati et al. (2020) addressed a multi-allocation variant of the HLRP with time windows in which both the hubs and vehicles are assumed to be capacitated and distance balancing is done by a bi-objective framework. In their model, pick-up and delivery were not done simultaneously, but for both actions, the routes had the same visiting sequence. They proposed a multi-objective imperialist competitive algorithm (MOICA), which was compared with a non-dominated sorting genetic algorithm (NSGA-II). The MOICA turned out to find better solutions for large scaled instances, requiring over 1200 seconds for instances of 100 nodes.

3.4.2.4 Local Search-Based Approaches

Local Search (meta)heuristics aim at exploring the solution space by moving the current solution to another promising solution in the neighbourhood (Mekamcha et al., 2021). Rieck et al. (2014) presented a fix-and-optimize heuristic for their variant of the many-to-many location routing problem from the timber-trade industry which is suitable for large-scale instances. In the fix-and-optimize scheme, in every iteration, there is a subset of binary variables whose values are fixed in order to limit the dimension of the branch-and-bound tree and speed op the computational time.

Lopes et al. (2016) proposed a variant of the MMLRP in which routes contain exactly one hub each and there is one additional route connecting all hubs, this variant is called the many-to-many p-location hamiltonian cycle problem (MMpLHP). In their article, three different solution approaches are explored: a multi-start variable neighbourhood descent (VND), a biased random key genetic algorithm (BRKGA), and a local solver. They experimented with different settings and exploration strategies for the VND and found that it performed better than the genetic algorithm as it provided better solutions over various test instances in less computational time.

Next to the PSO, Ghaffarinasab et al. (2018) proposed an iterative weiszfeld-type algorithm (IWA). This IWA was able to provide equally high-quality solutions as the PSO, yet for the 10 polygon areas, it only took almost 18 seconds.

Karimi & Setak (2018) proposed a bi-objective formulation for flow shipment in an incomplete network, meaning it is not interconnected. They proposed a MILP formulation and a normalized weighted-sum method (NWS) and an ε-constrained method to solve well-known instances and a case study on the Iran Post company. Their method allowed them to lower computation time by 89.34% compared to their base model for the case study. Solving the 30-node problem in less than 1.5 hours.

Abbasi et al. (2019) compared a variable neighbourhood search (VNS) approach against a benders decomposition and variable neighbourhood particle swarm optimization methods on the many-to-many HLRP as proposed by de Camargo et al. (2013). They tested 35 sample tests ranging from 10 to 1000 nodes with a maximum run time of 6 hours. They found that the VNS was more efficient than the VNPSO in samples of up to 440 nodes, and with samples of up to 1000 nodes only the VNS was able to find a solution within the set time boundaries (5.9h for 1000 nodes). Also in maritime logistics, the HLRP can be applied.

Fontes & Concalves (2021) proposed an HLRP formulation using sub-hubs, which are intersection points in two nearby regions allowing demand to be transported by local ships without the use of additional hubs or deep-sea services. To solve their formulation, they use a cutting plane approach for small, and a variable neighbourhood decomposition search (VNDS) for large instances. The VNDS was able to solve instances of up to 200 nodes in around 20 minutes.

Another variant of the HLRP is introduced by Real et al. (2021), who propose the multimodal hub network design problem with flexible routes (MHNDPFR). Flexible routes mean that there may be a mix of hub and non-hub nodes in a route. They propose a top-down and bottom-up ALNS heuristic (TDALNS; BUALNS) and apply those to benchmark instances of up to 50 nodes. The advantage of MHNDPFRs over classical models is that no particular topological structure such as cycles, stars or trees is assumed a priori (Real et al., 2021) allowing for the most cost-efficient routes.

Ratli et al. (2022) built further on the findings of Lopes et al. and proposed a general variable neighbourhood search (GVNS). The GVNS employs a VND as a local search but does not use a completely random re-start to resolve local optimum traps as is the case with the variants proposed by Lopes et al. In their GVNS, the search sequentially explores seven different neighbourhood structures instead of only one. The results of the GVNS are good as it was able to improve the best-known solutions of 691 out of 912 instances and was able to solve large instances fast (about 40 times faster than the VND approaches proposed by Lopes et al. (2016)).

Lastly, an adaptive large neighbourhood decomposition search (ALNDS) meta-heuristic was proposed by Wu et al. (2022) to solve both the multi-allocation HLRP (MAHLRP) and single-allocation HLRP (SAHLRP). They applied both heuristics on instances of the Australian Post dataset and showed that they are more efficient than CPLEX and that multi-allocation can efficiently reduce operating costs compared to single allocation. Depending on the 'tightness' of the problem formulation, computation times for 70 node instances varied strongly between 20 and 100 minutes.

3.4.2.5 Hyper Heuristic Based Approaches

Danach et al. (2019) proposed a hybrid hyper-heuristic approach for the capacitated single allocation p-hub location routing problem (CSApHLRP). The approach is hybrid since the hyper heuristic benefits from a relax-and-cut method, more specifically the information obtained by solving the dual problem through Lagrangian relaxation. The low-level heuristics are guided by a learning method (reinforcement learning) known as association rules, this technique originates from data mining and aims to find relationships between different implemented heuristics to find the best series of heuristics to be applied. The largest tested instances contained 100 nodes and 10 hubs, which were solved in over 4 minutes.

Pandiri & Singh (2021) addressed the many-to-many p-location hamiltonian cycle problem (MMpLHP) variant, earlier described by Lopes et al. (2016). They have proposed two hyper-heuristic approaches, based on random and greedy selection mechanisms, that acquired much better performance in terms of solution quality as well as execution time compared to the BKRGA, VND, and LS proposed by Lopes et al..

3.4.2.6 Synthesis

Below a synthesis of the main aspects, based on the overview given by Yang et al. (2019), of the works of the abovementioned authors is given in Table 5. Within this table, 'Problem Type' gives the classification of the problem by the author(s). 'Hub Capacity' indicates whether hub capacity is taken into account by the author(s). 'Number of hubs' indicates whether the number of hubs is fixed (p-hubs) or unfixed, Non-Hub Allocation indicates whether the author(s) use single- or multiple-allocation, 'Routing Constraints' indicate the type of routing constraints used by the author(s), 'Solution Method' indicates the chosen solution approach, 'Application/Data' indicates the dataset that the solution approach is applied to, 'Problem Size' indicates the (largest) number of nodes that the solution approach is applied on, 'Pick-up/Delivery' indicates whether or not pick-up and delivery actions are allowed in one route. Furthermore, although it is not included, Real et al. (2021) consider a heterogenous fleet in contrast to other authors.

As can be observed in Table 5, there are many different formulations and many different solution approaches for the HLRP. We see that most authors use a single allocation strategy and do not consider hub capacity. Furthermore, the routing constraints can be one out of four types: length, capacitated, time, or the number of nodes. In terms of the problem size the solution is applied to, we see mostly problem sizes of up to 200 nodes, which is not the size of (our) real-life cases. Except for a few exact approaches most solutions are population-based followed by local searches. However, for the larger instance problems, we see that only local search-based solutions are used. Lastly, the article that relates most to our problem is that of Wasner & Zäpfel (2004), followed by Seto et al. (2021) because of their time window formulations. In Table 4, we provide some abbreviations that are used in Table 5 but were not previously mentioned.

Table 4 Abbreviations not earlier introduced

AP	Australian Post standard data set	**OR-LIB**	Library for test data sets for OR problems
CAB	Civil Aeronautics Board data set	**BDA**	Benders Decomposition Algorithm
TSPLIB	Sample library for TSP (related) problems	**URAND**	Random data set
TR	Turkish network data set		

Table 5 An overview of literature on the Hub Location Routing Problem

Authors	Problem Type	Hub Capacity	Number of Hubs	Non-Hub Allocation	Routing Constraints	Solution Method	Application/ Data	Problem Size	Pick-up/ Delivery
Nagy & Salhi, 1998	MMLRP	Yes	Unfixed	Single	Length + Capacitated	Hierarchical	One instance	249	Simultaneous
Wasner & Zäpfel, 2004	HLVRM	Yes	Unfixed	Multi + Direct	Length + Capacitated	Hierarchical	Austrian Parcel	2042	Distinct
Çetiner et al., 2010	HLRP	No	p Hubs	Multi	Length	Two-Stage Heuristic	Turkish Postal	81	Simultaneous
Catanzaro et al., 2011	PHLRP	No	Unfixed	Multi	Num. Nodes	Branch-and-Cut	Random Instances	20	Distinct
de Camargo et al., 2013	MMHLRP	No	Unfixed	Single	Time	BDA	AP	100	Simultaneous
Mokhatari & Abbasi, 2014	MMHLRP	No	Unfixed	Single	Time	VNPSO	Random Instances	300	Simultaneous
Rodríguez-Martín et al., 2014	HLRP	No	p Hubs	Single	Num. Nodes	Branch-and-Cut	AP & CAB	50	Distinct
Rieck et al., 2014	MMLRP	No	p Hubs	Single + Direct	Capacitated	Multi-Start + GA	Timber trade Industry	140	Distinct & Simultaneous
Bostel et al., 2015	HLRP	Yes	Unfixed	Single	Num. Nodes	CPLEX; MA	AP postal	100	Simultaneous
Lopes et al., 2016	MMpLHP	No	p Hubs	Single	Num. Nodes	VND; BRKGA; LS	76 instances from TSPLIB	14 to 783	Simultaneous
Kartal et al., 2017	pHLRP-SPD	No	P Hubs	Single	Num. Nodes	MSA; ACO	CAB, TR, AP, URAND	400	Simultaneous
Ghaffarinasab et al., 2018	HLRP	No	p Hubs	Single	Capacitated	IWA; PSO	-	-	Distinct
Karimi & Setak, 2018	(S)BO-FSS-IHLRP	No	Unfixed	Single	Time	NWS; ε-Constr.	AP; Iran Postal	30	Distinct
Abbasi et al., 2019	MMHLRP	No	Unfixed	Single	Time	VNS	Random Instances; OR-LIBRARY	1000	Simultaneous
Danach et al., 2019	CSApHLRP	No	Bounded	Single	Capacitated	Hyper Heuristic	AP	100	-
Kartal et al., 2019	pHCVRP	No	p Hubs	Single	Num. vehicles	ACS, DPSO	TR, AP	200	Distinct
Yang et al., 2019	HLRP	Yes	Unfixed	Single	Capacitated	CPLEX; MA	AP	100	Distinct
Basirati et al., 2020	MMHLRP	Yes	Unfixed	Multi	Capacitated + Time + length	MOICA; NSGA-II	CAB	100	Simultaneous
Fontes & Goncalves, 2021	MMpHLRP(SH)	No	p Hubs	Single/Multi	None	Cutting Plane; VNDS	AP	200	-
Pandiri & Singh, 2021	MMpLHP	No	p Hubs	Single	Num. Nodes	HH_RAND; HH_GREED	76 instances from TSPLIB	14 to 783	Simultaneous
Real et al., 2021	MHNDPFR	No	Unfixed	Multi	Time + Capacitated	TDALNS; BUALNS	Random; AP; CAB	50	-
Seto et al., 2021	MMHLRP	No	Unfixed	Multi	Time + Capacitated + Num. Nodes	Heuristic	Random Instances	50	Distinct
Ratli et al., 2022	pHLRP	No	p Hubs	Single	Num. Nodes	GVNS	76 instances from TSPLIB	14 to 783	-
Wu et al., 2022	MMHLRP	Yes	Unfixed	Multi	Capacitated	ALNDS	AP	50/70	Simultaneous
This Research	MMHLRP	No	Unfixed	Single	Time + Capacitated		Dutch Parcel	4000	Distinct

3.5 SOLUTION APPROACHES

In this section, we address the most interesting solution approaches from the literature (Table 5). Many authors use different approaches, which all have their characteristics, advantages, and disadvantages. Furthermore, they are difficult to compare due to the variety of application instances. Some authors use the same classical datasets like CAB or AP. Yet those datasets often are small and are not comparable to real-life instances.

Given the size of our problem of over 4000 locations, two main characteristics are important for determining the suitability of a solution approach. The first is the size of the problem to which the approach has been applied since that is an indication that it is also suitable for our problem. The second is the computation time required to solve the largest instances the solution approach has been applied to, as we expect an algorithm that is fast on 'smaller' sized problems will also perform relatively fast on large-sized problems.

Below we will discuss different types of approximate solution methods and their suitability. Exact methods are not discussed since they do not apply to the size of the problem. Furthermore, not all methods found in the literature will be discussed. Discussed methods are either relevant given the two criteria given above or because they have been applied by multiple authors.

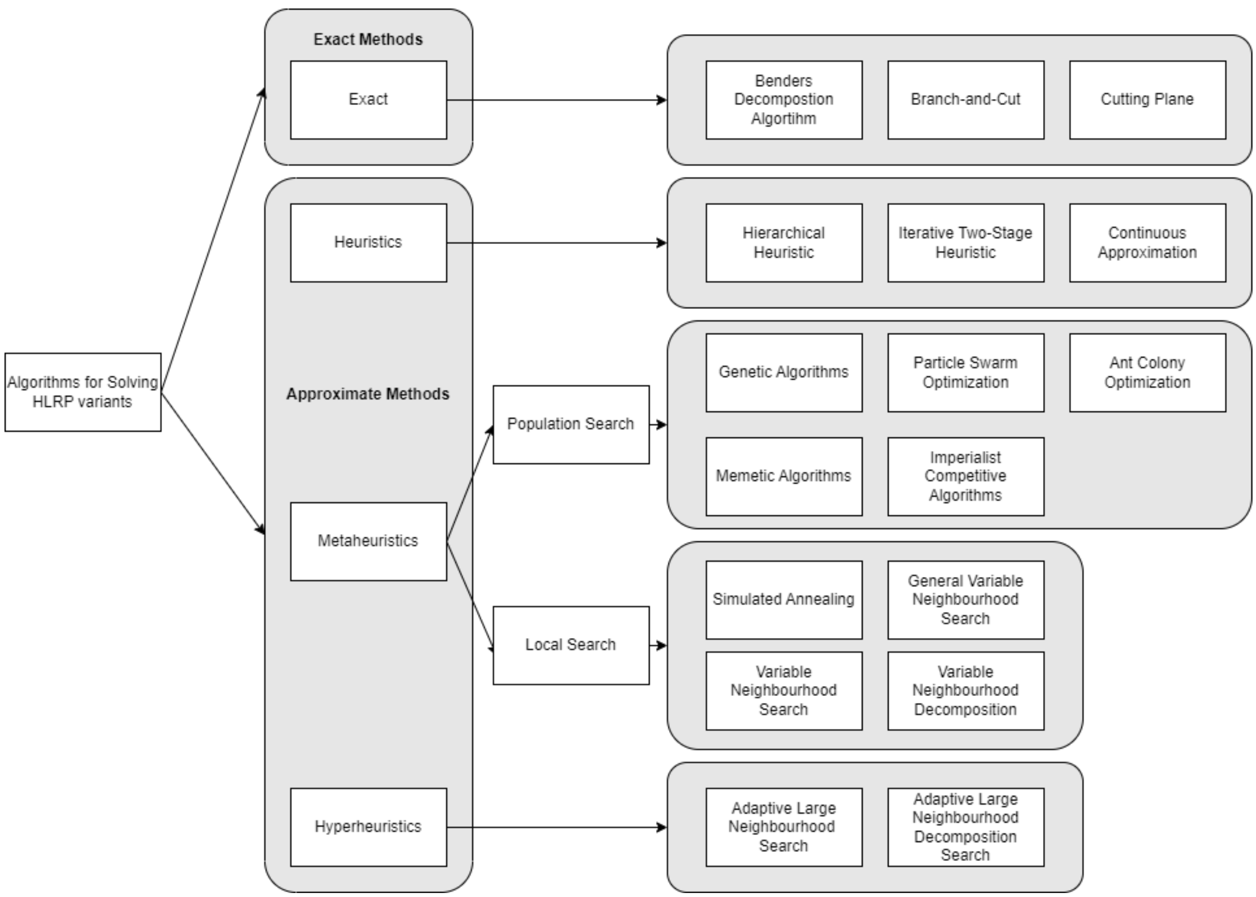

Figure 14 Overview of solution methods for the Hub Location Routing Problem

3.5.1 Heuristics

Multiple authors use some kind of heuristic approach, although mostly as part of a larger meta- or hyper-heuristic. One heuristic that is worth investigating is that of Wasner & Zäpfel. Already in 2004, they applied a heuristic on an Austrian Postal Service company in a network of 2042 postal zones. This is the largest problem we have found in literature and therefore relevant to our research. Other heuristics as the Two-stage heuristic

(Çetiner et al., 2010) or continuous approximation (Ghaffarinasab et al., 2018) are not investigated due to the small problem instance or unrelatedness.

In their heuristic, Wasner & Zäpfel build their solution sequentially and the results from one problem area form the constraints for the next. By means of the feedback loops, the solution is improved iteratively. Although the results of the method are promising (14.7% cost reduction), there is no information on the computation time that is required. Therefore, it is hard to assess the suitability of this method from a tractability perspective. Therefore, we can conclude that the sequential building of a solution is an effective approach (hence great cost reduction), yet it remains unknown whether this method is a practical one.

3.5.2 Metaheuristics

Metaheuristics are search methods that balance intensification and diversification to overcome local optima. Many different types of metaheuristics have been applied to the HLRP. We discuss variable neighbourhood search and general variable neighbourhood search as they have been applied to large problems (1000 nodes) or showed low computation times on relatively large problems. Furthermore, we discuss adaptive large neighbourhood (decomposition) search. Although often applied, we will not address population search metaheuristics as they have neither been applied to large problems nor showed low computation times.

3.5.2.1 Variable neighbourhood (decomposition) search

In a variable neighbourhood search (VNS), each iteration improvements are explored by applying different neighbourhood structures to one solution. Abbasi et al. (2019), applied a variable neighbourhood search on a large problem instance of 1000 nodes, which required over 21,000 seconds to solve.

In the variable neighbourhood decomposition search proposed by Fontes & Goncalves (2021), the problem is decomposed into subproblems to which either a local search or a variable neighbourhood descent (VND) is applied. VND is a simplified VNS in which one takes an initial solution, keeps improving it within a neighbourhood until no better solutions can be found, and then applies a different neighbourhood following a predetermined sequence. This process repeats a number of k times (Lopes et al., 2016). After a number of k_{max} iterations, the best solution thus far is destroyed and then the process repeats a fixed number of times. This method achieved solutions for instances of 200 nodes in around 1250 seconds. Its advantage is that it thoroughly explores different neighbourhood structures and decomposes the problem into subproblems (which speeds up the process). Yet, it has a stopping criterion a fixed number of iterations, which could create the situation that the algorithm terminates while it is still exploring promising solutions.

3.5.2.2 General variable neighbourhood search

The application of a general variable neighbourhood search (GVNS) to the HLRP is proposed by Ratli et al. (2022), which is a variant of the variable neighbourhood search. Their GVNS also employs VND as a local search but does not use a completely random re-start to resolve local optimum traps. Instead, it modifies the current solution to varying degrees to resolve a local optimum trap. Furthermore, it applies a set of different neighbourhood structures in a fixed sequential way, which differs from other VND approaches. All in all, their algorithm creates an initial solution, changes that solution, and then it applies a VND of seven sequential neighbourhood structures, saving the best-found solution and repeating this process until k_{max} iterations. The algorithm is terminated when the maximum running time has been met.

This approach has been applied on known test datasets on which it showed very good performance, improving 691 out of the 912 known best solutions while being capped at 60 seconds. In these 912 problems, some are larger-sized instances of a few hundred nodes. The advantage of this approach is that it seems to attain good solutions in a short time, implying it is still tractable for very large problem instances. Furthermore, it is relatively simple to implement when the neighbourhood structures and sequence are determined. The exploration however is always random, and with large instances, an adaptive exploration strategy might be needed.

3.5.2.3 Adaptive large neighbourhood search

A large neighbourhood search (LNS) is a local search framework in which solutions are relaxed (elements are removed from the solution) and re-optimized (these elements are placed back) leading to better solutions. This method was introduced by Shaw (1998). Following the LNS framework, Ropke & Pisinger (2006) proposed the adaptive large neighbourhood search (ALNS). In this meta-heuristic framework solutions are destroyed and repaired by making use of multiple neighbourhood operators. These operators are chosen based on their weights, and these weights are determined based on their performance. Thus, when a better solution has been found, the used operators' weights are increased by a factor or decreased when a worse solution has been found.

Real et al. (2021) proposed a framework in which they implemented two ALNS approaches which were applied sequentially within one iteration, one which started with determining hub locations (top-down ALNS) and one which started with constructing routes (bottom-up ALNS). Although this is an interesting approach, it required relatively high computation times (over 1500 seconds) for small instances (up to 50 nodes). Therefore, it is not deemed suitable enough for our problem.

3.5.2.4 Adaptive large neighbourhood decomposition search

As an expansion to the ALNS, Wu et al. (2022) propose an adaptive large neighbourhood decomposition search (ALNDS). In this framework, the destroy and repair phase is not applied to the whole solution space but only to a reduced solution space that corresponds to a subproblem of the original problem. For both the destroy and repair phases operators are used consecutively and are chosen based on their weights, which follow from the combined weight of that operator to be chosen as well as the weight of the subproblem at hand. Selection is handled by a roulette wheel mechanism. An interesting addition that Wu et al. made is that to facilitate better exploration, they allow non-feasible solutions to be accepted while compensating with adaptive weights.

The ALNDS was applied to a limited set of nodes (70) and required relatively much computation time to solve (6100 seconds). Therefore, although the adaptive nature of the ALNDS seems very useful, we consider it not feasible to apply ALNDS in the real-life case of Company X.

3.5.3 Hyper heuristics

A hyper heuristic is a heuristic method that selects or generates lower-level heuristics. In HLRP literature, some hyper heuristics are applied, of which some showed good performance.

The first hyper-heuristic is that of Danach et al. (2019), which works similarly to an ALNS. They created a method that uses online learning (or reinforcement learning) to iteratively determine the best series of lower-level heuristics to apply in sequence. The weights of the lower-level heuristics are determined by their average objective value of the times the heuristic is applied. The hyper heuristic can choose from a set of constructive, improvement and perturbation heuristics. Although the problem size to which they applied the hyper heuristic is not very large (100 nodes) the required time to solve the largest instances is also relatively small (240 seconds). Since it is a learning algorithm, we expect that its efficiency becomes higher when applied to larger problems.

Lastly, Pandiri & Singh (2021) also applied a hyper-heuristic. They used two heuristics, one that has a random selection mechanism and one that uses a greedy selection mechanism. They applied their heuristics on larger instances than Danach et al. and performed well compared to methods proposed by Lopes et al. (2016). The computation time seems to be slightly better than that of Danach et al., yet it is growing large for larger problems. Therefore, the practicality (of any ALNS) for this problem is not entirely clear.

3.6 DECISION-MAKING IN REAL-LIFE APPLICATIONS

The hub location routing problem is, as indicated by Çetiner et al. (2010), a strategic or tactical problem. Yet, in the long(er) term parameters like, e.g., costs, demands, and distances tend to change or are harder to predict. Therefore, it is risky to make strategic decisions based on the result of solving one problem instance and can lead to sub-optimal decisions (Contreras, 2015; Snyder, 2006). Especially in the case of Company X, there is fluctuation in the daily workload as well as the workload per month. On top of that, the demand on the network will also change if the e-commerce company will increase its volumes, thus longer-term forecasts also need to be taken into account when making such strategic decisions. Therefore, we want to find out what is been described in the literature to use optimization algorithms in such strategic decision-making in real-life applications, i.e., dealing with uncertainty.

First of all, it has become clear that many authors that published on the HLRP use test data sets to compare the performance of their algorithm, and thus do not address how their algorithm can be used to make real-life strategic network design decisions (Abbasi et al., 2019; Bostel et al., 2015; Danach et al., 2019; Fontes & Goncalves, 2021; Gelareh et al., 2015; Ghaffarinasab et al., 2018; Kartal et al., 2017, 2019; Mokhatari & Abbasi, 2014; Pandiri & Singh, 2021; Ratli et al., 2022; Real et al., 2021; Rodríguez-Martín et al., 2014; Seto et al., 2021; Wu et al., 2022; Yang et al., 2019). Furthermore, some authors do use their proposed algorithm to solve real-life cases, yet do not specify their exact approach or only use/present one solution (Basirati et al., 2020; Catanzaro et al., 2011).

Then some authors in HLRP literature do touch upon the topic of using algorithms for decision-making in real applications. Wasner & Zäpfel (2004) state that due to the size of the HLRP, practical cases are only solvable with problem-specific knowledge, i.e., simplification, restriction or partial linearization of the problem. Furthermore, they do not indicate the use of different scenarios in determining the optimal network structure. This implies that they use only one, representative, dataset. Çetiner et al. (2010) describe how they apply different parameter settings to find the cheapest network configuration. However, since they indicate that they lack detailed mail flow data we can only assume that have used only one input dataset on which the network is optimized. To deal with variable cost components that differ depending on a short or long planning term, Rieck et al. (2014) chose to create one representative planning period of a week. Lastly, in their concluding remarks, Yang et al. (2019) indicate that decision-making for real applications of the HLRP does not reduce to experimenting on a single data set but involves evaluating and comparing a significant number of alternative scenarios and performing sensitivity analysis on significant parameters as well as assessing the robustness of the determined hub locations. Furthermore, having to assess many scenarios requires an efficient solution technique that can find realistic solutions in a reasonable time, devising such a technique was out of their scope.

3.6.1 Stochastic & Robust Optimization

The HLRP literature is not decisive on how to deal with decision-making in real-life applications, therefore we expand our scope into HLP, LRP, and facility location literature, but first, we will elaborate on decision-making. Rosenhead et al. (1972) divide decision-making into three categories:

1. **Certainty**, all parameters are deterministic and known.
2. **Risk**, there are uncertain parameters whose values are governed by known probability distributions. These problems are known as stochastic optimization problems, which common goal is to optimize the expected value of the objective function (Snyder, 2006).
3. **Uncertainty**, situations and parameters are unknown as well as their probability distributions. These problems are known as robust optimization problems and often attempt to optimize the worst-case performance (minimax) of the system (Snyder, 2006).

The goal of both stochastic and robust optimization is to find a solution that will perform well under any possible realization of the random parameters, what “performing well” means differs per application (Snyder, 2006). Random parameters can either be continuous or described by discrete scenarios. This scenario approach has two main drawbacks, firstly creating good scenarios is difficult, secondly for computational reasons, one wants to identify a relatively small number of scenarios which limits the range of future states under which the decisions are evaluated (Snyder, 2006). However, the scenario approach generally results in more tractable models. For stochastic optimization, the most common objective is to minimize the expected cost. For robust optimization, the two most common measures are minimax cost and minimax regret, which can often be transformed into each other. Regret is the difference between the cost of a solution in a given scenario and the cost of the optimal solution for that scenario (Snyder, 2006).

For dealing with uncertainty in a hub location problem, Alumur et al. (2012) consider two sources of uncertainty. Namely, the set-up costs for hubs and demands to be transported between the hubs. To deal with the first source, they propose a minimax regret model to minimize the worst-case over a finite set of scenarios. To deal with the second source, they propose a two-stage stochastic linear program with recourse. To deal with both sources simultaneously, they propose a stochastic minimax regret MIP formulation with recourse. A similar approach was done by Habibzadeh Boukani et al. (2016), although they used robust optimization, instead of stochastic programming, in a minimax formulation over five different scenarios. Yahyaei et al. (2014) used a design of experiments (DOE) to create many different scenarios that each was optimized on three different performance indicators. Lastly, we want to mention more authors applying robust optimization to the HLP (Khaleghi & Eydi, 2022; Shahabi & Unnikrishnan, 2014), or facility location problem (Gülpınar et al., 2013).

To solve a capacitated multi-echelon location-routing problem (quite similar to the HLRP), Winkenbach et al. (2016) use multiple techniques in order to make their solution suitable for real-life applications. Firstly, they split their algorithm into two independent subproblems, and secondly by using approximations for routing costs instead of using explicit routing. They did not include the so-called 'robust optimization' in their method. They do, however, indicate the practical and academic relevance of doing so, especially for the robust optimization by making use of 'uncertainty sets' so that the optimal solution works well not only for a particular scenario but rather for a broad set of scenarios.

All in all, we want to know how to make valid strategic decisions in real-life applications by making use of an, in principle, deterministic algorithm. We found that in the current HLRP literature, there is not much said about how to determine good future-proof solutions. What was said was the necessity to limit the problem size as much as possible and use either one representative dataset or use multiple scenarios to find good solutions. Expanding into HLP, LRP, and facility location literature we found that there are different categories of decision-making under uncertainty. Depending on the category, we found that either stochastic programming or robust optimization is used to find more robust solutions.

3.6.2 Simheuristics

Another method to deal with large complex combinatorial optimization problems (COPs) in an uncertain (stochastic) environment is making use of simheuristics. Simheuristics allow modellers for dealing with real-life uncertainty in a natural way by integrating simulation into a metaheuristic-driven framework (Juan et al., 2015). Metaheuristics benefit from different random-search and parallelization paradigms, but they often assume the parameters, objective and constraints to be deterministic (Chica et al., 2020; Juan et al., 2015). Simulation can be understood as the process of model 'execution' that takes a model through its evolution over time (Chica et al., 2020).

Simheuristics allow for modelling uncertainty by making sure that the feedback from the simulation is used to guide the metaheuristic search process itself, and all the information obtained by the simulation component allows for considering a risk/reliability analysis on stochastic solutions to the stochastic optimization problem (Chica et al., 2020). Many authors use monte carlo simulation with their metaheuristic framework (Chica et al., 2020). A general scheme of a simheuristic can be seen in Figure 15. This general simheuristics approach described by Juan et al. (2015) has two main characteristics:

1. It promotes a closer integration between optimization and simulation. In particular, the evaluation of solutions is performed not only by simulation but also by problem-specific analytical expressions. Hence, it mixes simulation and ad hoc approximations making them appealing for optimization, but they do not represent reality accurately.
2. The feedback from simulation can be used not only to evaluate solutions but also to refine the analytical part so that the latter can generate and/or evaluate more realistic solutions

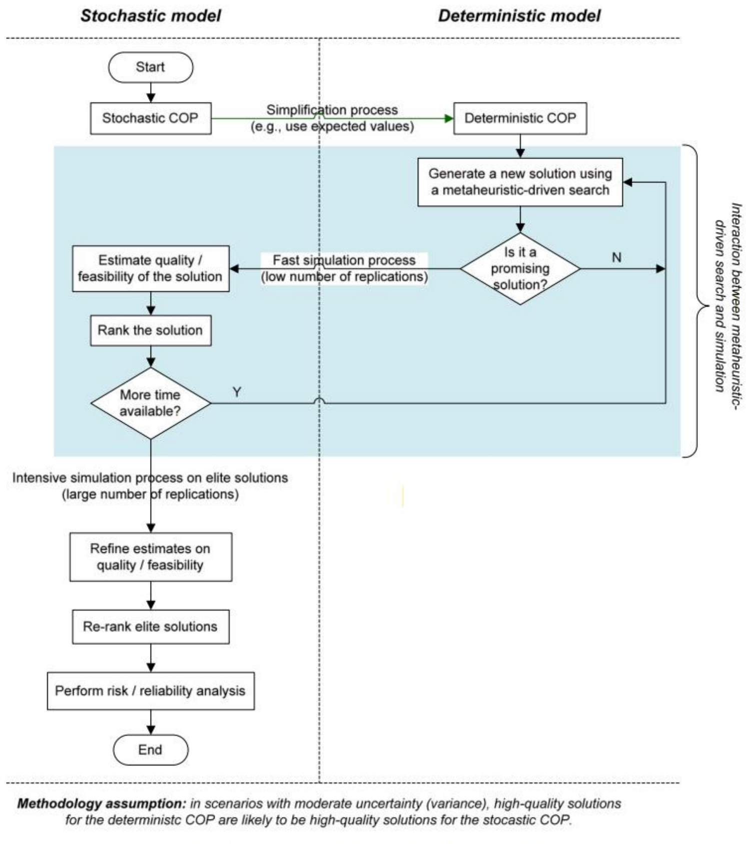

Figure 15 General scheme of simheuristics for solving stochastic COPs (Juan et al., 2015)

The use of simheuristics comes with advantages and disadvantages. Based on the reviews of Juan et al. (2015) and Chica et al. (2020) the following advantages of using simheuristics can be identified:

- Simheuristics allow decision-makers to construct and study valid models of complex systems that produce high-quality solutions in reasonable computing times.
- Simheuristics allow for risk assessment of alternative solutions and sensitivity analysis.
- Simheuristics allow for better system understanding and output analysis.

Simultaneously, they also described some limitations of simheuristics:

- Results are not expected to be truly provably optimal. Metaheuristics do not ensure to result in an optimal solution, this is amplified by using a simulation to be optimized.
- An effort from additional stakeholders is demanded to define the system. By making use of a 'white-box' paradigm, i.e. being transparent, external stakeholders are required to validate the simulation system.
- More computational resources are required compared to traditional methods because of the integration of a simulation engine with a metaheuristic.
- An additional trade-off in diversification versus intensification has to be made in the solutions that are sent to the simulation component, as the uncertainty might lead to worse solutions being selected, and the number of replications of the simulation as the metaheuristic proceeds.

Lastly, Chica et al. (2020) also determined some best practices for the design and implementation of simheuristics:

- Do not overload simheuristics with long simulations. Use strategies to limit the use of simulation time so that the metaheuristic has time to converge. For example, by using a three-stage approach in which the simulation intensity is increased with the stages, each consists of less but higher quality solutions to evaluate.

- Choose a simulation paradigm that is understandable to decision-makers.
- Choose an appropriate simulation paradigm for each stage of a simheuristic. For example, use more computationally expensive paradigms (e.g., agent-based simulation) for the last stages while using lighter computational paradigms (e.g., Monte Carlo simulation) for the first stages.
- Validate the simulation model before running the simheuristic. Make sure that the simulation is an accurate representation of the real-world system.

Two interesting applications of simheuristics are Calvet et al. (2019), who have applied a simheuristic approach to the stochastic multi-depot VRP for which they could relax strong assumptions that previous works had to apply to solve the problem, and de Armas et al. (2017), who applied a simheuristic to the uncapacitated facility location problem. Both used an iterated local search (ILS) metaheuristic combined with monte carlo simulation.

3.7 CONCLUSIONS ON LITERATURE

In this chapter, we researched the literature to get a good understanding of network design, specifically the design of hub networks. We found that there are multiple optimization problems concerned with hub network design. Given the goal of finding a network design that minimizes the operational costs, we can conclude that the HLP and LRP do not provide all the required elements. We found that the class of hub location routing problems, which is concerned with the optimization of hub locations combined with routing costs, is the most suitable for our problem. Because the HLRP combines the strategic hub locations, allocation of service areas, and the routing it encompasses all the elements of our problem. Due to the complexity of these problems, this class is relatively new and has not been researched much. We have reviewed articles by 24 authors and examined the problem characteristics they used, and the solution approaches they applied.

In the literature, we have seen many different solution methods to solve (variants of) the HLRP. Based on two criteria, the size to which the methods are applied and the computation time required to solve a problem instance, we have found three types of solution methods that might be suitable. The first is the heuristic approach with feedback loops of Wasner & Zäpfel since it has the most similar application. The second is the (general) variable neighbourhood search that has been applied to a large problem of 1000 nodes and Ratli et al. and has shown very good performance in terms of computation time.

Therefore, we conclude that the best solution approach for the case of Company X will be to apply a general variable neighbourhood search. Additionally, we looked into approaches to solve optimization problems for real life in which there is uncertainty to deal with. We found that there are three categories of uncertainty, namely certainty, risk, and uncertainty. A distinguishing factor between the latter two types if whether the uncertainty can be modelled using probability distributions. Additionally, we discussed simheuristics, which allow modellers for dealing with real-life uncertainty in a natural way by integrating simulation into a metaheuristic-driven framework. This method is very suitable for our problem since, due to its strategic nature, we have to deal with a largely deterministic dataset which in the simheuristic can be used to model stochasticity.

4 MODELLING THE OPTIMAL NETWORK

In this chapter, we will answer the third research question '*How can we model and solve the logistics network of Company X?'* by formulating a mathematical model that can solve the network of Company X. From the literature in the previous chapter, we know which solution methods are most suitable for solving the model and how to deal with stochasticity in the problem. Therefore, in this chapter, we will formulate a solution method that can solve the logistical network and can deal with the stochastic demand. In this chapter we will answer the following sub-questions:

- What are the assumptions needed to model the network?
- What is the scope of the model?
- How can we solve the logistical network of Company X?

4.1 ASSUMPTIONS AND APPROXIMATION ON POSTAL ZONES

Because of the size and complexity of the problem, we need to make simplifying assumptions and use approximations instead of explicit calculations to keep the model tractable.

4.1.1 Assumptions

In the modelling of the HLRP for Company X, we make the following assumptions to simplify the model:

1. We assume the backbone network, meaning the inter-hub connections, to be fully interconnected.
2. We approximate the linehaul cost by multiplying the linehaul volume times a certain discount factor per parcel per kilometre $0 \leq \alpha \leq 1$.
3. There is a finite distinct set of possible hub locations.
4. Each hub must be associated with at least one local route.
5. Every tour must start and end at the hub it is associated with.
6. Collection and delivery are distinct processes and cannot happen simultaneously in the same route.
7. There is no limit on the capacity of each hub.
8. There is no limit on the number of vehicles that can be used.
9. There is a minimum and maximum number of hubs that can be opened.

4.1.2 Postal zones

In order to keep the size of the model small, we need to aggregate the volumes and number of stops on a PC4 postal zone level, leading to a bound on the number of stops. In the Netherlands, PC4 postal zones are the areas with addresses with the same numbers in their postal code.

Within the Netherlands there are 4070 postal zones, providing an upper bound on the number of locations to consider in the model. The aggregation process is a little different per process (PPC, BPC, HD), for example during PPC there is a distinction between pickup retailers and deposit retailers.

For the PPC process, the first step will be to determine per postal zone how many retailers there are to service. The second step is to determine how many pickup retailers and deposit retailers a postal zone contains. Thirdly, we determine per postal zone for how much volume those two groups of retailers account. For the BPC and HD processes, aggregating is similar but there is no distinction to make. Therefore, we can determine per postal zone the total volume that needs to be collected/delivered and the number of visits that need to be made.

Using the volumes, we can determine how many 'postal zones' fit in a route. Using both the volumes and number of stops, a total service time of a postal zone can be determined which can be used to model how many postal zones fit in the process' time window.

4.2 HLRP DEFINITION OF THE NETWORK OF COMPANY X

Let $G = (V, E)$ be an undirected complete graph with vertices set V and edge set E, with $E = \{(i, j): i, j \in V, i \neq j\}$, see Figure 16 Example of G, without edges

. Let $H \subset V$ be a subset of V containing all the vertices that can function as a hub node. Let R, D, B, and C all be subsets of V containing all vertices with pickup retailers, deposit retailers, batch retailers, and consumers respectively. Because each vertex represents a postal zone, it can contain multiple types of retailers as well as consumers and can therefore be in more than one of the aforementioned sets. Thus, V is the union of all unique items in the sets H, R, D, B, and C, i.e., $V = H \cup \{R \cup D \cup B \cup C\}$.

Figure 16 Example of G, without edges

With each node $h \in H$ a fixed cost of FCH_h is associated with opening a hub in that node. Furthermore, with each node $i \in \{R \cup D \cup B \cup C\}$ a volume and service/stop time is associated. On top of that, every node i must be allocated to a hub node. With every edge $e \in E$ a driving distance and time cost, D_{ij} and TT_{ij} respectively, is associated with, $D_{ij}, TT_{ij} \in \mathbb{R}^{+}$.

A tour T is a subgraph of G such that nodes $V[T]$ form a sequence of distinct nodes $\{h, v_i, \dots, v_j, h\}$ and edges $(i, j) \in E$ that starts and ends with at a certain hub node h. With each tour, a driving range and timespan are associated, DR_t and DT_t respectively.

For each process, a graph C must be created which is the union of a set of opened hub nodes and tours such that all nodes in the sets of vertices corresponding to that process ($\{R \cup D\}, B, C$) are allocated to a hub and a tour. In the creation of these subgraphs, constraints like vehicle capacity, driving range, and time windows need to be satisfied.

Since the network of Company X must deal with different types of logistical processes, we define a possible solution S as a combined graph consisting of a set of hub nodes $\{h_h, h_k, \dots, h_m\} \in H$ and three subgraphs C_1, C_2, C_3 that represent the tours (and thus service area) for each process. Examples of subgraphs C_1, C_2, and C_3 can be seen in Figure 17 An example of subgraph C_1, the PPC process, Figure 18 An example of subgraph C_2, the BPC process, Figure 19 An example of subgraph C_3, the HD process, and .

The goal is to find the solution graph S that minimizes the total network cost. The total network costs consist of the cost of opening hubs, the linehaul between depots, the locating of collection points, usage of vehicles, and the cost of traversing the edges for each logistical process.

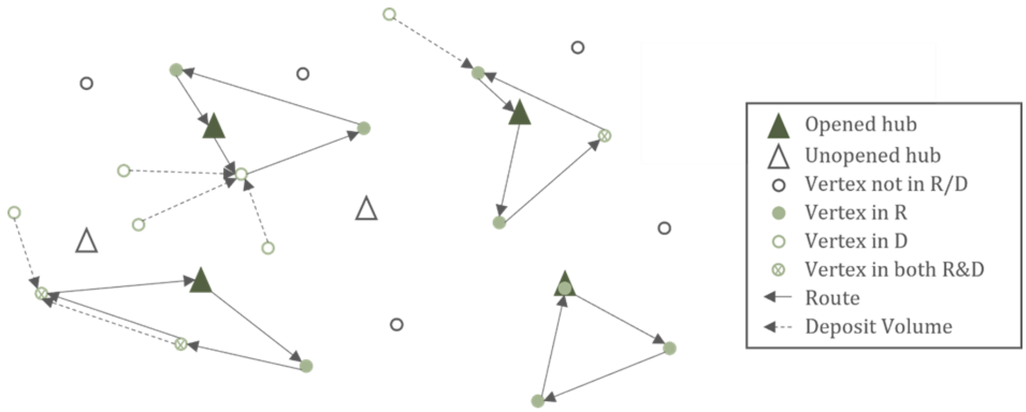

Figure 17 An example of subgraph C_1, the PPC process

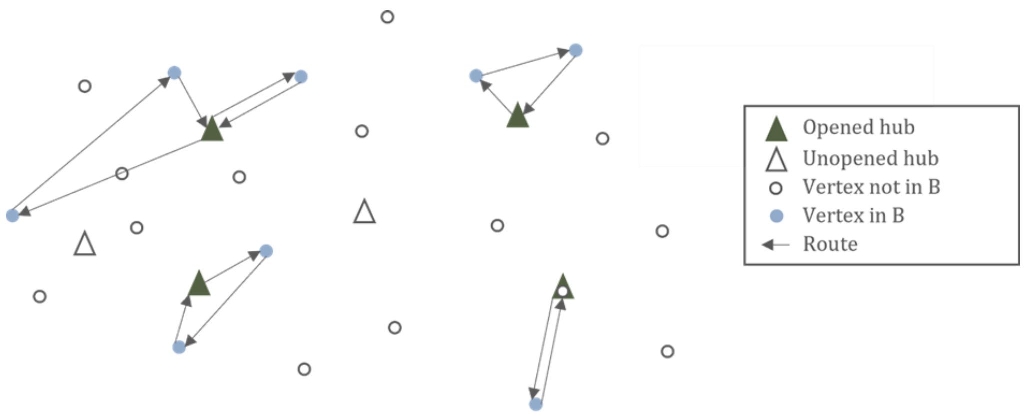

Figure 18 An example of subgraph C_2, the BPC process

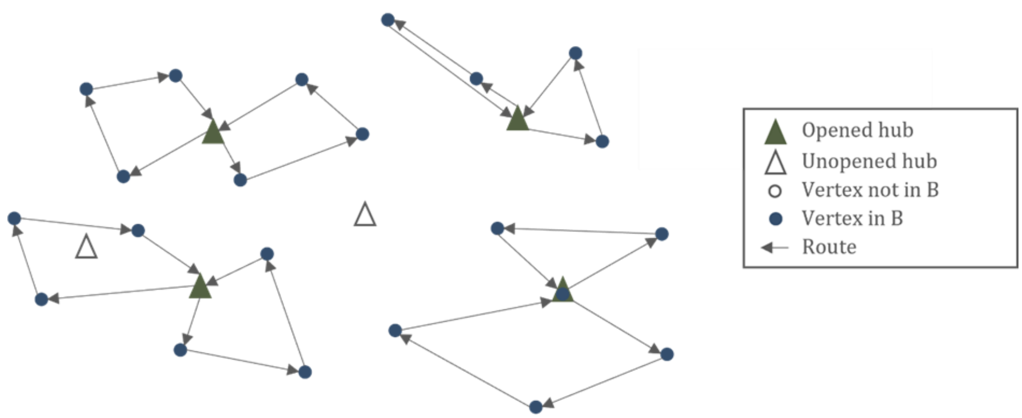

Figure 19 An example of subgraph C_3, the HD process

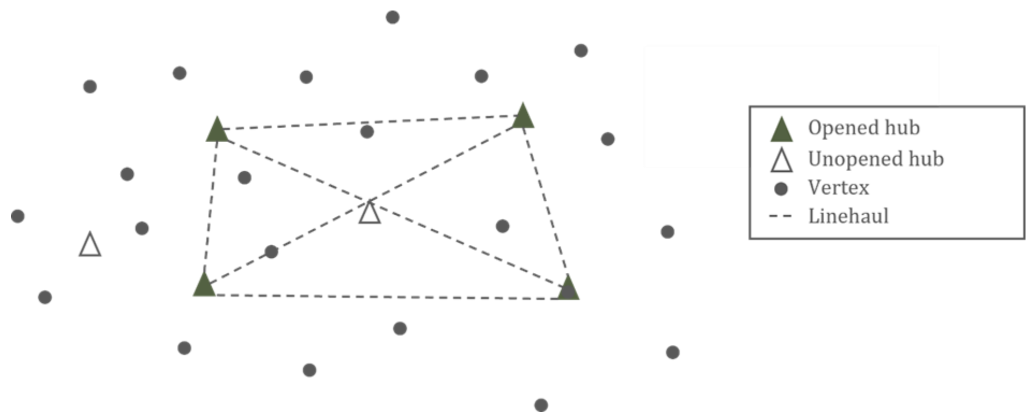

Figure 20 An example of the resulting linehaul of the subgraphs C_1 and C_3

4.3 MIXED-INTEGER LINEAR PROGRAMMING MODEL

In this section, we provide and explain the model sets, parameters, variables, and constraints. Like most authors in HLRP literature, we start by formulating a mixed integer linear programming model (MILP). This MILP formulation is inspired by the HLRP formulations of multiple authors, specifically, the time window constraints are inspired by Seto et al. (2021). We point out that we have formulated the model such that it represents reality as accurately as possible. This means that the routing in the heuristic, which will be an approximation, will be different than the routing in the MILP.

4.3.1 Sets

H Set of all possible hub postal zones
R Set of all pick-up retailer postal zones
B Set of all batch parcel retailer postal zones
D Set of all deposit retailer postal zones
V Set of all vehicles
C Set of all delivery postal zones
P Set of all processes (PPC, BPC, HD)
N Combined set of all retailer postal zones ($\{B \cup R \cup D\}$)
A Combined set of all postal zones ($\{H \cup R \cup D \cup B \cup C\}$)

4.3.2 Parameters

FCH_h Fixed Cost of opening hub h
VQ_v The capacity of vehicle v
TCV_v Cost per hour of travelling with vehicle v
VCV_v Cost per km of vehicle v
FCV_v Fixed cost of using vehicle v
D_{ij} Distance between postal zone i and postal zone j
TT_{ij} Travel time to travel from postal zone i to postal zone j $(i,j) \in \{H \cup R \cup D \cup C\}$
PQ_{rc} Pick-up Quantity of postal zone r meant for postal zone c during PPC
DOQ_{dc} Deposit Quantity in postal zone d destinated for postal zone c
BQ_b Pick-up Quantity of postal zone b during BPC
DQ_{ic} Delivery Quantity at postal zone c originating from postal zone i
TP_i Required time for serving postal zone i
TA_i^p First possible time of arrival at location i for process p
SV_v^p Suitability of vehicle v for process p, $SV_v^p \in \{0,1\}$
DR_v Driving Range of vehicle v
DT Maximum distance between a deposit postal zone and the collection point zone

α Line-haul cost per unit per km

4.3.3 Variables

ho_h Variable indicating if hub h is opened yes or no, $ho_h \in \{0,1\}$
uv_v^{hp} Variable indicating if vehicle v is used by hub h during process p, $uv_v^{hp} \in \{0,1\}$
qv_v^p Quantity of volume loaded into vehicle v during process p
ha_i^{hp} Whether or not location i is assigned to hub h
π_{ikmj} The total volume that from postal zone i to postal zone j via hub k and m
πr_{ikmj} The pickup volume that travels from postal zone i to postal zone j via hub k and m
πd_{ikmj} The deposit volume that travels from postal zone i to postal zone j via hub k and m
do_d^i Variable indicating whether or not location $d \in D$ deposits its volume at location $i \in \{R \cup D\}$
dov_d^v Auxiliary variable indicating whether the deposit volume of retailer d is picked up by vehicle v
cp_i Whether postal zone i contains a collection point or not
x_{ij}^{phv} Whether vehicle v, which departed from hub h, is travelling from postal zone i to postal zone j during a route in process p, $x_{ij}^{hpv} \in \{0,1\}$
ta_i^{pv} The arrival time of vehicle v for process p at location $i \in A$
td_i^{pv} The departure time of vehicle v for process p at location $i \in A$
pv_i^{pv} Volume picked up at location $i \in N$ by vehicle v during process p
hac_{ij}^{km} Auxiliary variable indicating if the connection i-k-m-j exists
β_i^{pv} fraction of volume of location i that is picked up by vehicle v during process p

4.3.4 Model

Objective:

$$\min z = \sum_{h \in H} ho_h \cdot FCH_h + \sum_{i \in N} \sum_{j \in C} \sum_{(k,m) \in H} \alpha \cdot \pi_{ikmj} \cdot D_{km} + \sum_{v \in V} \sum_{h \in H} uv_v^h \cdot FCV_v + \sum_{v \in V} \sum_{h \in H} \sum_{p \in P} \sum_{(i,j) \in A} x_{ij}^{hpv} \cdot (VCV_v \cdot D_{ij} + T_{ij} \cdot TCV_v)$$

Subject to

Hub, Vehicle, and Collection Point allocation:

$$ha_i^{hp} \leq ho_h \quad \forall\, i \in N, \forall\, p \in P, \forall\, h \in H \tag{1}$$
$$do_d^i \leq cp_i \; \forall\, i \in N, \forall\, d \in D \tag{2}$$
$$uv_v^{hp} \leq ho_h \quad \forall\, v \in V, h \in H, \forall\, p \in P \tag{3}$$
$$\sum_{i \in N} do_d^i + cp_d = 1 \quad \forall\, d \in D \tag{4}$$
$$\sum_{h \in H} ha_i^{hPPC} \geq 1 \quad \forall\, i \in R \tag{5}$$
$$\sum_{h \in H} ha_d^{hPPC} \geq cp_d \quad \forall\, d \in D \tag{6}$$
$$ha_d^{hPPC} + 1 \geq ha_i^{hPPC} + do_d^i \quad \forall\, d \in D, \forall\, i \in R, \forall\, h \in H \tag{7}$$
$$\sum_{h \in H} ha_i^{hBPC} = 1 \quad \forall\, i \in B \tag{8}$$
$$\sum_{h \in H} ha_c^{hHD} = 1 \; \forall\, c \in C \tag{9}$$
$$do_d^i \cdot TT_{d,i} \leq DT \quad \forall\, i \in \{R \cup D\}, \forall\, d \in D \tag{10}$$

Routing constraints:

$$\sum_{j \in A} x_{ji}^{hpv} = \sum_{j \in A} x_{ij}^{hpv} \; \forall\, i \in A, \forall\, v \in V, \forall\, p \in P, \forall\, h \in H \tag{11}$$

$$\sum_{j \in H} \sum_{i \in N} x_{ij}^{hpv} \leq 1 \quad \forall\, v \in V, \forall\, p \in P, \forall\, h \in H \tag{12}$$

$$\sum_{j \in H} \sum_{i \in N} x_{ji}^{hpv} \leq 1 \quad \forall\, v \in V, \forall\, p \in P, \forall\, h \in H \tag{13}$$

$$x_{ii}^{hpv} = 0 \quad \forall\, i \in A, \forall\, v \in V, \forall\, p \in P, \forall\, h \in H \tag{14}$$

$$x_{ij}^{hpv} = 1 - x_{ji}^{hpv} \quad \forall\, i \in N, \forall\, j \in A, \forall\, v \in V, \forall\, p \in P, \forall\, h \in H \tag{15}$$

$$x_{ij}^{hpv} \leq ha_j^h \;\; \forall\, j \in N, \qquad \forall\, i \in A, \forall\, p \in P, \forall\, v \in V, \forall\, h \in H \tag{16}$$

$$\sum_{i \in A} \sum_{h \in H} \sum_{v \in V} x_{ij}^{hPPCv} = 1 \quad \forall\, j \in R \tag{17}$$

$$\sum_{i \in A} \sum_{h \in H} \sum_{v \in V} x_{id}^{hPPCv} = cp_d \quad \forall\, d \in D \tag{18}$$

$$\sum_{i \in A} \sum_{h \in H} \sum_{v \in V} x_{ij}^{hBPCv} = 1 \quad \forall\, j \in \tag{19}$$

$$\sum_{i \in \{H \cup C\}} \sum_{h \in H} \sum_{v \in V} x_{ic}^{hHDv} = 1 \quad \forall\, c \in C \tag{20}$$

$$x_{ij}^{hpv} \leq uv_v^{hp} \quad \forall\, (i,j) \in A, \forall\, v \in V, \forall\, p \in P, \forall\, h \in H \tag{21}$$

$$\sum_{j \in N} x_{hj}^{hpv} = uv_v^{hp} \quad \forall\, v \in V, \forall\, h \in H, \forall\, p \in P \tag{22}$$

$$x_{ij}^{hpv} \leq ho_h \quad \forall\, i \in A, \forall\, j \in H, \forall\, v \in V, \forall\, p \in P, \forall\, h \in H \tag{23}$$

$$\sum_{h \in H} uv_v^{hp} \leq SV_v^p \quad \forall\, v \in V, \forall\, p \in P \tag{24}$$

$$\sum_{(i,j) \in N} \sum_{h \in H} x_{ij}^{hpv} \cdot D_{ij} \leq DR_v \quad \forall\, v \in V, \forall\, p \in P \tag{25}$$

Time window constraints:

$$td_i^{pv} + TT_{ij} \leq ta_j^{pv} + M \cdot \left(1 - x_{ij}^{hpv}\right) \;\; \forall\, (i,j) \in A, p \in P, v \in V, h \in H \tag{26}$$

$$ta_i^{pv} + TP_i \leq td_i^{pv} + 24 \cdot \left(1 - \sum_{j \in A} x_{ji}^{phv}\right) \;\; \forall\, (i,j) \in A, p \in P, v \in V, h \in H \tag{27}$$

$$ta_i^{pv} \geq TA_i^p \quad \forall\, i \in R, \forall\, p \in P, \forall\, v \in V \tag{28}$$

$$ta_d^{pv} \geq TA_d^p \cdot cp_d \quad \forall\, d \in D, \forall\, p \in P, \forall\, v \in V \tag{29}$$

$$ta_h^{pv} \leq TA_h^p \quad \forall\, h \in H, \forall\, v \in V, \forall\, p \in P \tag{30}$$

Capacity constraints:

$$pv_i^{PPCv} = \sum_{j \in A} \sum_{h \in H} x_{ji}^{hPPCv} \cdot \sum_{c \in C} PQ_{ic} + DOQ_{ic} \cdot cp_i + \sum_{d \in D} \sum_{c \in C} DOQ_{dj} \cdot dov_d^v \quad \forall\, i \in N, \forall\, v \in V \tag{31}$$

$$\sum_{i,j \in N} \sum_{h \in H} x_{ij}^{hPPCv} + \sum_{i \in N} do_d^i \leq 1 + dov_d^v \quad \forall\, i \in N, \forall\, d \in D, \forall\, v \in V \tag{32}$$

$$dov_d^v \leq \sum_{i,j \in N} \sum_{h \in H} x_{ij}^{hPPCv} \quad \forall\, v \in V, \forall\, d \in D \tag{33}$$

$$dov_d^v \leq \sum_{i \in N} do_d^i \quad \forall\, v \in V, \forall\, d \in D \tag{34}$$

$$pv_i^{BPCv} = \sum_{j \in A} \sum_{h \in H} x_{ji}^{BPChv} \cdot \sum_{j \in N \setminus \{i\}} PQ_{ij} \quad \forall\, i \in N, \forall\, v \in V \tag{35}$$

$$pv_i^{HDv} = \sum_{j \in A} \sum_{h \in H} x_{ji}^{HDhv} \cdot \sum_{j \in N \setminus \{i\}} PQ_{ji} \quad \forall\, i \in N, \forall\, v \in V \tag{36}$$

$$qv_v^p = \sum_{i \in N} pv_i^{pv} \quad \forall\, v \in V, \forall\, p \in P \tag{37}$$

$$qv_v^p \leq VQ_v \quad \forall\, v \in V, \forall\, p \in P \tag{38}$$

Linehaul constraints:

$$\pi r_{ikmj} = PQ_{ij} \cdot hac_{ij}^{km} \quad \forall\, i \in R, \forall j \in C, \forall\, (k,m) \in H \tag{39}$$

$$\pi d_{ikmj} = DOQ_{ij} \cdot hac_{ij}^{km} \quad \forall\, i \in D, \forall j \in C, \forall\, (k,m) \in H \tag{40}$$

$$\pi_{ikmj} = \pi r_{ikmj} + \pi d_{ikmj} \quad \forall\, i \in N, \forall j \in C, \forall\, (k,m) \in H \tag{41}$$

$$ha_i^k + ha_j^m \leq 1 + hac_{ij}^{km} \quad \forall\, i \in N, \forall\, j \in C, \forall\, (k,m) \in H \tag{42}$$

$$hac_{ij}^{km} \leq ha_i^{kPPC} \quad \forall\, i \in N, \forall\, j \in C, \forall\, (k,m) \in H \tag{43}$$

$$hac_{ij}^{km} \leq ha_j^{mHD} \quad \forall\, i \in N, \forall\, j \in C\; \forall\, (k,m) \in H \tag{44}$$

$$\sum_{(k,m) \in H} \sum_{(i,j) \in N} \pi_{ikmj} = \sum_{(i,j) \in N} PQ_{ij} + DOQ_{ij} \tag{45}$$

Sign constraints:

$$ho_h \in \{0,1\} \quad \forall\, h \in H$$
$$uv_v^h \in \{0,1\} \;\; \forall\, v \in V, \forall\, h \in H$$
$$ha_i^{hp} \in \{0,1\} \;\; \forall\, i \in N, \forall\, p \in P, \forall\, h \in H$$
$$do_j^i \in \{0,1\} \;\; \forall\, i,j \in N$$
$$cp_i \in \{0,1\} \quad \forall\, i \in N$$
$$dov_j^v \in \{0,1\} \;\; \forall\, j \in N, \forall\, V \in V$$
$$x_{ij}^{phv} \in \{0,1\} \quad \forall\, i,j \in N, \forall\, p \in P, \forall\, h \in H, \forall\, v \in V$$
$$qv_v^p \geq 0 \;\forall\, v \in V, p \in P$$
$$\pi_{ijkm} \geq 0 \;\forall\, i,j \in N, \forall\, k,m \in H$$
$$ta_i^{pv} \geq 0 \;\forall\, i \in N, \forall\, p \in P, \forall\, v \in V$$
$$td_i^{pv} \geq 0 \;\forall\, i \in N, \forall\, p \in P, \forall\, v \in V$$
$$pv_i^{pv} \geq 0 \;\forall\, i \in N, \forall\, p \in P, \forall\, v \in V$$

4.3.5 Objective and constraints explanation

In this section, the objective function and all 45 constraints are explained. First the objective function, it represents the total cost of the network which consists of:

- fixed cost of opening a hub;
- The linehaul costs;
- The fixed cost for using vehicles;
- Routing costs (both travelled distance and time travelled).

4.3.5.1 Hub & Collection Point allocation

Constraint (1) makes sure that a location can only be allocated to a hub if that hub is opened. Constraint (2) allows a deposit retailer d only to deposit its volume in postal zone i if this zone contains a collection point. Constraint (3) prohibits vehicles to be allocated to unopened hubs. Constraint (4) makes sure that a deposit retailer either drops its volume or is a collection point itself. Constraints (5)–(9) indicate for each process that the postal zones that need to be visited are allocated to a hub. Here, constraint (7) makes sure that a deposit retailer in postal zone d is allocated to the same hub as the postal zone to which it is allocated. Constraint (10) indicates that the postal zone i to which the deposit retailers in postal zone d deposit their products cannot be more than DT hours away.

4.3.5.2 Routing constraints

Constraint (11) is a balance flow constraint. Constraints (12) and (13) indicate that each vehicle can only enter and leave a hub once. Constraint (14) makes sure that a vehicle cannot travel to itself. Constraint (15) makes sure that a vehicle cannot travel back to the location that it came from. Constraint (16) makes sure that a location cannot be visited from hub h if it is not allocated to hub h. Constraints (17) – (20) make sure that every postal zone that needs to be visited is visited. Constraint (21) makes sure that a postal zone can only be reached by a used vehicle. Constraint (22) makes sure that each used vehicle departs from the hub it is allocated to. Constraint (23) makes sure that a movement by vehicle v that is used by hub h can only be done if that hub is opened. Constraint (24) makes sure that for each process vehicles can only be used when they are suitable for that process and simultaneously makes sure that per process a vehicle can only be allocated to one hub. Lastly, constraint (25) makes sure that a vehicle cannot drive more than its maximum driving range. Because, each vehicle must depart from a depot, can only enter and leave a node once, and cannot travel back to the node it came from, no sub-tours can exist.

4.3.5.3 Time window constraints

Constraints (26) and (27) make sure that when during process p vehicle v drives to a postal zone the arrival time is equal to the departure time of the previous postal zone plus the travel time and that the departure time at a postal zone is equal to its arrival time plus its service time. Constraints (28) – (30) make sure that vehicles arrive after the start of the time window at postal zones and are back at the hub before the end of the time window.

4.3.5.4 Capacity constraints

The capacity constraints, (31) – (38), make sure that the volumes going into vehicles are calculated correctly. Constraint (31) makes sure that the volume picked by vehicle v at postal zone i during PPC is equal to the own pickup retailer's volume of that location plus the volume of all deposit retailers that deposit their parcels there (indicated by dov) and the volume of zone i's own deposit retailers. Constraints (32) – (34) form an AND-constraint indicating that volume deposited by retailers in postal zone d at a collection point in postal zone i is transported by vehicle v if both d is connected to collection point i and zone i is visited by vehicle v. Constraint (35) and (36) calculate the volumes to be collected and delivered at batch parcel retailers and consumers respectively. Constraint (37) calculates the volume loaded in vehicle v during process p and constraint (38) makes sure that this volume does not exceed the capacity.

4.3.5.5 Linehaul constraints

Constraint (39) calculates the pickup volume that travels from postal zone i via hubs k and m to postal zone j by multiplying the pickup volume at i destined for j times the variable indicating whether or not the connection between i-k-m-j exists Constraint (40) does the same but for the deposit volume of postal zone i which can be routed differently due to the collection point allocation. The total volume that travels from postal zone i to postal zone j is determined by constraint (41). Constraints (42) – (44) form an AND-constraint indicating that the connection i-k-m-j only exists if zone i is allocated to hub k during PPC and location j is allocated to hub m during HD. Constraint (45) makes sure that the total linehaul volumes are equal to the total volumes to be shipped. Constraint (44) is similar to (42), only does this constraint represent that the connection d-k-m-j between a postal zone d's deposit retailers and a customer j can only exist if deposit retailers in d deposits their volumes at a collection point in postal zone i and zone i is allocated to hub k during PPC and postal zone j is allocated to hub m during HD. Constraint (45) makes sure that the total of all volumes shipped between postal zones i and j is equal to the volume from zone i destined for zone j.

4.3.6 Testing the model

To make sure that the mathematical model is correct, we have tested it for a very small instance with three possible hub locations and ten postal zones represented as nodes in a two-dimensional plane. To keep the model simple, only the PPC (0) and HD (1) processes are used. Nodes 0-6 contain pickup retailers, and thus are in set R. Nodes 7-9 contain only deposit retailers and are thus in set D. All postal zones have consumers to whom parcels must be delivered and thus are in set C. There are four tours available for PPC, and three for HD, which are smaller. The resulting network, see Figure 21 Solution of the MILP model for a very small test instance of 10 nodes

, opens one hub, from which one tour visits all the nodes in R. All nodes in D deposit their volumes at other nodes, indicated by the orange, green, and red lines. For the HD process, two tours are required to visit all the nodes. The required run time to solve this small instance is 2404 seconds.

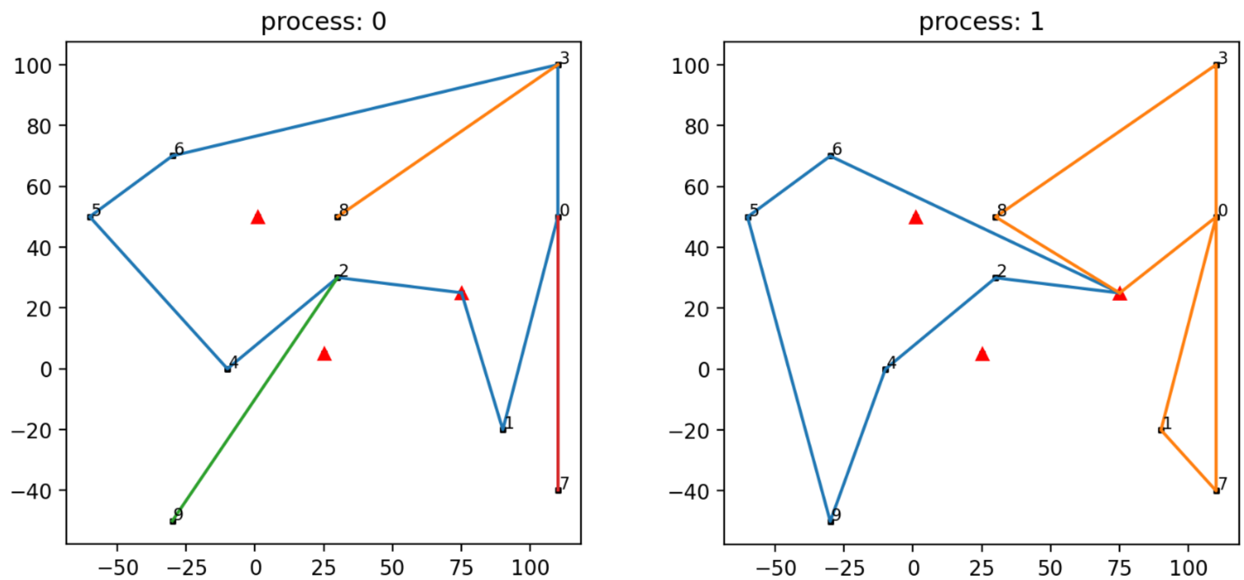

Figure 21 Solution of the MILP model for a very small test instance of 10 nodes

The routes are as follows:

- Process 0: hub-1-0-3-6-5-4-2-hub
- Process 1: hub-2-4-9-5-6-hub
- Process 1: hub-8-3-7-1-0-hub

4.4 GVNS-BASED HEURISTIC

The number of variables required to model the network grows very large, with theoretically every postal zone being a possible hub location the number of variables can grow to 4070^4 = $O(10^{12})$. Since the model is a MILP model, solving it becomes intractable very quickly. Even more so with the number of variables in the model, that we need to handle. Therefore, we devised a heuristic approach to find high-quality solutions within a reasonable time.

From the literature, we know that there is a wide variety of heuristic algorithms that have been applied to HLRPs. For our problem, we found that the most applicable methods are the heuristic with feedback loops of Wasner & Zäpfel (2004) due to the similar application case and the general variable neighbourhood search (GVNS) of Ratli et al. (2022). In our heuristic approach, we have combined these two methods.

In this section, we will address the creation procedure of the initial solution, the used operators, and the separate procedures that are used in the GVNS.

4.4.1 Creating an initial solution

The creation of an initial solution consists of two steps, the initial hub opening and allocation and the collection point opening and allocation.

4.4.1.1 Hub opening and allocation

For the creation of the initial solution, we will consider three different procedures. The quality of each procedure is evaluated in experiment 3 in Section 5.3.3.

The first procedure is a random procedure. It starts by choosing a number of hubs to open. This number is a random number, x, between the minimum and the maximum number of hubs that are allowed to be opened, these values are given as input to the procedure. Consecutively, randomly x hubs are chosen to be opened. Then for each postal zone, it is determined how long the driving time is to each opened hub and the hub that has the least driving time is chosen. This is repeated until each postal zone is allocated.

The second procedure is a greedy clustering approach, see lines 1-16 of Algorithm 1. The procedure takes in a predefined number of hubs to open k. If k is not passed, it is chosen to be the ceiling of the average of the minimum and maximum number of hubs that are allowed to be opened. Then, the service area of each hub is determined by taking the set of postal zones that have a drive time of up to 60 minutes to the hub. While the number of opened hubs is less than k, the hub is chosen that adds the highest number of postal zones to the set of zones that are already serviced (i.e., are in the service area of an earlier opened hub). If no hubs add new postal zones, the hub is chosen that has the best average drive time to the hubs in its service area. After a new hub has been determined, that hub is opened and the list of serviced postal zones is updated. When k hubs are opened, all postal zones are assigned to the hub to which they have the lowest driving time and the costs (demand-weighted driving time) are calculated.

The third procedure is a continuation of the greedy clustering approach and is described in Algorithm 1. It starts by determining the clusters in a greedy manner. Subsequently, it will try to improve the hub locations. This is done by testing whether moving a hub to a neighbouring hub location yields an overall improvement. Because a single hub location change might not make the overall solution better while multiple changes might yield an improvement, we introduced a diversification variable $divers$. This variable is initialized at 1, thus we start by trying to move the first hub that is in the $openHubs$ set to the location of its closest neighbour. If this yields an improvement, we save the new best set of opened hubs and costs and reset the variables for diversification and the number of unsuccessful tries. The more iterations in which no better solution is found, the more hubs are replaced to other locations. When there have been k iterations in which no improvement is found, the algorithm terminates and the best hub locations are used as initial solution.

Algorithm 1: Initial Solution

```
Function Initial Solution (k):
if k = ∅ then:
        k = ⌈(minNumHubs + maxNumHubs) / 2⌉
for h = 1 to |hubs| do:
        serviceArea_h ← Determine_ServiceArea(h, time = 60)
end
serviced, openHubs = ∅
while |openHubs| < k do:
        newZones ← DetermineNewZones(serviced)
        if |newZones| > 0 then:
                addHub ← argmax_{h∈H}(newZones_h)
        else:
                addHub ← argmin_{h∈H}(avgTime_h)
        serviced ← serviced ∪ newZones_addHub
        openHubs ← openHubs ∪ {addHub}
repeat
AssignPostalZonesToClosestHub( )
bestHubs ← openHubs
bestCosts ← CalculateDemandWeightedTime( )

divers = 1
tries = 0
while tries ≤ k do:
        neighHubs = ComputeClosestNeighbourForEachHub( )
        for s = 1 to divers do:
                openHubs ← openHubs \ {openHubs_s}
                openHubs ← openHubs ∪ {neighHubs_s}
        end
        AssignPostalZonesToClosestHub( )
        costs = CalculateDemandWeightedTime( )
        if costs < bestCosts then:
                bestCosts ← costs
                bestHubs ← openHubs
                tries ← 0
                divers ← 1
        else:
                tries ← tries + 1
                divers ← divers + 1
repeat
```

4.4.1.2 Collection point opening and allocation

Next to the allocation of postal zones to hubs, there is also the allocation of postal zones to collection points. Although the improvement of collection point locations is out of scope for this research, we need to model them to get accurate results. Therefore, we have created a greedy approach to locating and allocating to collection points.

The main idea is that ideally the collection points are located in postal zones that need to be visited anyway because large retailers are located there. Thus, as initialization, a list is created of all postal zones and the corresponding volume of pickup retailers that is sorted in descending order. Starting at the top of the list we check if the postal zone contains a collection point, whether it is allocated to a collection point, and whether it has a collection point within the maximum allowed range. If the postal zone is not allocated to a collection point, neither has one itself nor has none nearby, then a collection point is opened in the postal zone and it gets allocated to itself. If the postal zone is not allocated but has collection points within range, it gets allocated to the closest collection point.

Because we first evaluate the postal zones with pickup retailers, the number of additional stops due to collection points remains as low as possible and thus is a good approximation of the best collection point network.

4.4.2 Neighbourhood structures

In designing the logistical network there need to be made decisions on multiple levels: the hub level, the service area level, and the routing level. Therefore, for each level, we use different neighbourhood structures that represent the decisions that can be made on that level.

On the hub level, we have neighbourhood structures that together form the hub-level neighbourhood N_h:

- **add_hub**: randomly add a hub that is currently unopened and assign all the postal zones that lie within a range of $minTimeToHub$ minutes to the newly opened hub.
- **remove_hub:** randomly remove a hub that is currently opened and assign all the postal zones in the service area of that hub to the closest opened hub. A hub can only be closed if the number of opened hubs is above a certain threshold.
- **swap_hubs:** randomly choose one opened hub and one closed hub. Close the opened hub and open the previously closed hub. Assign the service area of the previously opened hub to the previously closed hub.

On the service area level, we have the neighbourhood structures that together form the service area level neighbourhood N_{sa}:

- **move_zone_PC2:** move a whole PC2 area from the service area of one hub to another. The choice for the PC2 area to move is made by taking either the PC2 area that has the largest average distance to the allocated hub, the largest average drive time to the allocated hub or randomly each with a probability of 1/3.
- **move_zone_PC3:** move a whole PC3 area from the service area of one hub to another. The choice for the PC2 area to move is made by taking either the PC3 area that has the largest average distance to the allocated hub, the largest average drive time to the allocated hub or randomly each with a probability of 1/3.
- **swap_zones:** choose two postal zones randomly and swap their hub allocations for all processes.
- **move_zone_Allp:** choose a postal zone and move it to the service area of one of the opened hubs for each of the processes. With a probability of 1/3 the chosen postal zone has the largest average distance to its allocated hub, with a probability of 1/3 the chosen postal zone has the largest average travel time to its allocated hub, and with 1/3 probability the choice is made randomly.
- **move_zone_1p:** choose a postal zone and move it to the service area of one of the opened hubs for one of the processes. With a probability of 1/3 the chosen postal zone has the largest distance to its allocated hub during that process, with a probability of 1/3 the chosen postal zone has the largest travel time to its allocated hub during that process, and with 1/3 probability the choice is made randomly.

Lastly, there is the routing level. To keep the model tractable, there is only one operator on this level which forms the routing neighbourhood N_{route}:

- **Determine_Routing:** this operator determines the total routing length, duration, and the number of routes per hub using the well-known savings heuristic of Clarke & Wright for the capacitated vehicle routing problem. In this algorithm, there is a single hub and a set of locations to visit. Every location will be connected directly to the hub. After which for each location pair (i, j) a saving is calculated by $S_{ij} = C_{i0} + C_{0j} - C_{ij}$, where 0 is the hub and C_{ij} is the cost of travelling from location i to location j. A savings list is created in which all location pairs (i, j) are sorted descending based on savings. The locations are added to routes starting at the top of the savings list until the list has been exhausted

(Clarke & Wright, 1964). An addition is that we included a penalty cost for each link that takes more than one hour to traverse. This penalty is included to avoid long distances in the allocation of postal zones that do not fit in other routes. These allocations can occur when single postal zones contain more volume than the capacity of a vehicle and are added as single-stop routes. This implementation is chosen to not have to add the complexity of dividing a single postal zone over multiple vehicles. The value chosen for the penalty is €250.

To spend the computation time to improve the service areas and not fix them, there is also an operator that assigns every postal zone to the closest hub. This operator can be used to reset the service areas after the hub locations have been 'shuffled' so that the service area operators can be used to improve those service areas.

4.4.3 General Variable Neighbourhood Search

The main algorithm is the general variable neighbourhood search (GVNS), see Algorithm 2. In this algorithm, both the shaking procedure and the variable neighbourhood descent procedure are called to find high-quality solutions. The function takes a solution, a k_{max}, T_{max}, and neighbourhoods N as input. The T_{max} puts a limit on the duration of the algorithm and such forms the stopping criterion. The parameter k_{max} denotes the maximum number of iterations that may be performed within a single shake. Choosing k_{max} is important as it determines the level of diversification of the GVNS. It is known that for high values of k_{max} the GVNS behaves like a random multi-start heuristic since the characteristics of the best-known solution get lost. Therefore, we want to keep the value of k_{max} as low as possible while still getting high-quality results (Ratli et al., 2022).

Algorithm 2: General Variable Neighbourhood Search

Function $GVNS(Sol,\ k_{max} = \{k^0_{max}, k^1_{max} k^2_{max}\}, T_{max} = \{T^0_{max}, T^1_{max}, T^2_{max}\},$
$N = \{N_0, N_1, N_2\})$:
$operators = \{N_h, N_{sa}, N_{route}\} = \{N_0, N_1, N_2\}$
$Sol \leftarrow Intial\_Solution(Sol)$
$T \leftarrow CPU\_Time(\,)$
$\boldsymbol{while}\ T \leq T^0_{max}\ \boldsymbol{do}$:
  $k \leftarrow 1$
  $\boldsymbol{while}\ k \leq k_{max}\ \boldsymbol{do}$:
    $k \leftarrow k + 1$
    $Sol' \leftarrow Shaking(Sol,\ k,\ \{N_0\})$
    $Sol'' \leftarrow Multi\_Level\_VND(Sol',\ operators)$
    $\boldsymbol{if}\ Sol''\, better\ than\ Sol\ \boldsymbol{then}$:
      $Sol \leftarrow Sol''$
      $k \leftarrow 1$
  repeat
  $T \leftarrow CPU\_Time(\,)$
repeat
return *Sol*

A graphical representation of the functioning of the GVNS over multiple levels is displayed in Figure 22 Graphical representation of working of GVNS algorithm on multiple levels.. The algorithm starts on the hub level, where the GVNS is started with the parameters and neighbourhood corresponding to the first level ($k^0_{max}, T^0_{max}, N_0$). Within this GVNS an initial solution is created (but only on the hub level). Then while the run time has not exceeded the set limit and the number shake variable k has not reached the maximum value, three steps are performed. The first is the shaking procedure, which creates a new solution, the more iterations no improved solution has been found the greater the shake. After the shake, a variable neighbourhood descent (VND) procedure is called. This procedure takes in the shaken solution Sol' which it improves by sequentially applying the neighbourhood structures for that level. After the VND has been finished, the resulting solution Sol'' is compared to the best solution up to that moment Sol. If Sol'' is better, it becomes Sol and the shake variable k is reset to 1, after which the process starts over. Only on the hub level, between the shaking and VND procedure, the operator is called that assigns each postal zone to the closest hub. This is done to prevent the service area level GVNS to improve very bad service areas, which takes a lot of time due to the routing calculation, allowing the solution to converge faster.

Please note that the inputs of k_{max}, T_{max}, and N are all sets. This is because the heuristic calls the same procedure, but on a lower level, to determine the best possible lower-level solution as input to itself. Therefore, the settings for the lower-level procedure calls must be passed to the highest-level GVNS. Each time a lower-level GVNS is called, the sets with settings that are passed are excluding their first element (element 0). Therefore, regardless of the level the GVNS is on, the first element of the sets k_{max}, T_{max}, and neighbourhoods N determine the settings for that level. The passing of the sets is done in the variable neighbourhood descent (Section 4.4.5).

The calling of lower-level GVNS procedures can go down until the routing level where the VND is replaced with the routing heuristic which is immediately returned to the higher service area level, so there is no improvement on the routing. The 'exchange' of solutions and best improvements between each level form the feedback loops like in the heuristic of Wasner & Zäpfel (2004).

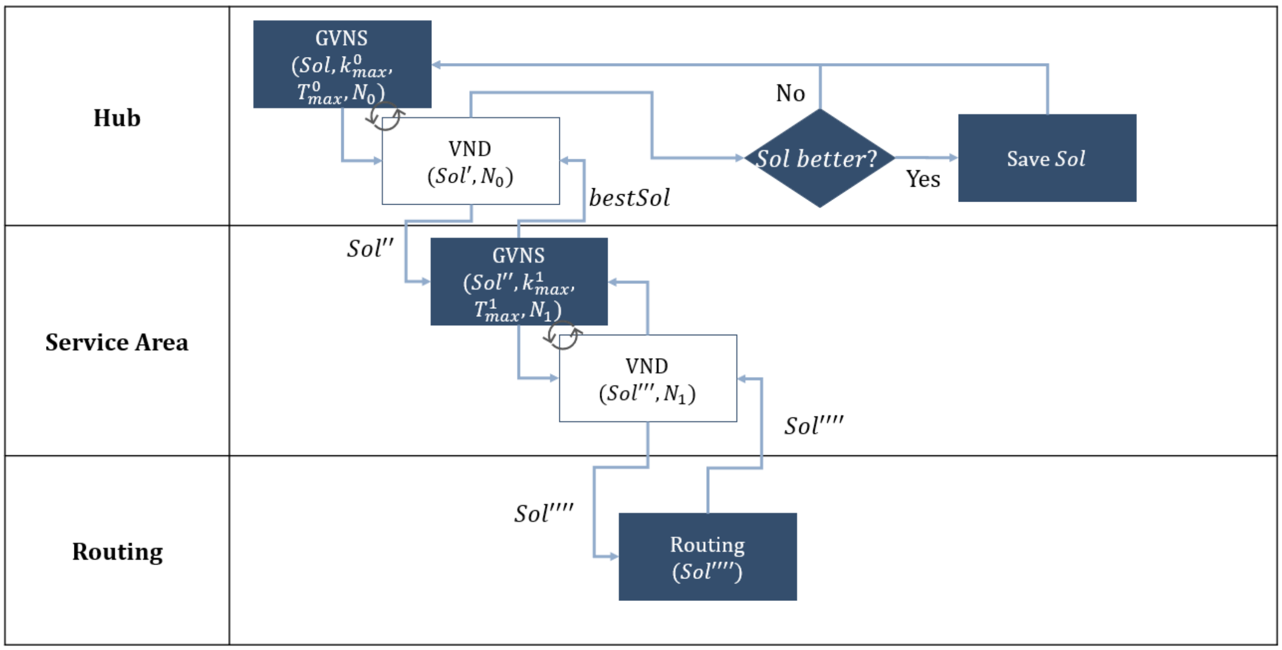

Figure 22 Graphical representation of working of GVNS algorithm on multiple levels.

4.4.4 Shaking Procedure

Like in Ratli et al. (2022), each iteration of the GVNS starts by calling a shaking procedure, see Algorithm 3. This procedure performs a fixed set of num_shake iterations in which it chooses an operator out of a given set of operators and modifies a solution consecutively and returns a changed solution. An import distinction from the variable neighbourhood descent is that this procedure aims to diversify and not to improve the solution.

Algorithm 3: Shaking Procedure

Function $Shaking\ (Sol, num\_shake,\ operators)$:
  $\boldsymbol{for}\ \ k = 1\ to\ num\_shake\ \boldsymbol{do}$:
    $Operator \leftarrow random.choice(operators)$
    $Sol \leftarrow Operator(Sol)$
  end
$\boldsymbol{return}\ Sol$

4.4.5 Multi-Level Variable Neighbourhood Descent

After the shaking procedure has been performed, the multi-level Variable Neighbourhood Descent (Algorithm 4) is called. This function is based on the Basic Sequential Variable Neighbourhood Solution that is used by Ratli et al. (2022) and is extended with the feedback loops as used by Wasner & Zäpfel (2004) to include the effect of lower-level performance in higher-level decisions. As input, the function takes a starting solution and a set of neighbourhoods as operators of which the last is always the *determine_routing* operator. It sequentially goes through the neighbourhood structures of the first neighbourhood in the set. While in higher-level neighbourhood structures (N_h, N_{sa}) it will also call a GVNS in the lower level to include the best possible performance on a lower level in the quality assessment of solutions at the higher level.

Algorithm 4: Multi-Level Variable Neighbourhood Descent

Function $Multi\_Level\_VND(Sol,\ operators, GVNS\_param)$:
$operators = \{N_h, N_{sa}, N_{route}\} = \{N_0, N_1, N_2\}$
$k \leftarrow 1$
    $\boldsymbol{while}\ k \leq |N_0|\ \boldsymbol{do}$:
        $Sol' \leftarrow N_0(k)(Sol)$
        $k \leftarrow k + 1$
        $\boldsymbol{If}\ |operators| > 1\ \boldsymbol{then}$:
            $Sol'' \leftarrow GVNS(Sol',\ N - \{N_0\},$
                $k_{max} - \{k^0_{max}\}, T_{max} - \{T^0_{max}\})$
            $\boldsymbol{if}\ Sol''\ better\ than\ Sol\ \boldsymbol{then}$:
                $Sol \leftarrow Sol''$
                $k \leftarrow 1$
        $\boldsymbol{else}$:
            $Sol \leftarrow Sol'$
    repeat
Return *Sol*

4.5 SIMHEURISTIC-BASED GVNS

Because a single deterministic scenario does not provide a good indication of whether a network design performs well in other scenarios. To deal with the stochastic nature of the problem, the deterministic GVNS will be expanded by making use of simulation.

The simheuristic expansion, as shown in Algorithm 5, requires some additional lines to the deterministic GVNS in Algorithm 2. The first is line 2, here an empty set is initialised which will be used to save all promising solutions that the GVNS will come across in. The next added line is line 10, this line makes sure that each time an improved solution is found in the deterministic GVNS it is tested in a quick simulation for a small number of scenarios to have some indication of its stochastic performance. In line 13, this shortly simulated solution is added to the list of promising solutions and saved for when the GVNS is finished. Lastly, line 15 contains the intensive simulation procedure which takes the m best solutions from the list of promising solutions and evaluates them in an intensive simulation with a high number of scenarios. The final result will be a list of the best-performing solutions and some statistical analysis of their performance.

Note that the GVNS with simulation is only called on the highest level. The Multi_Level_VND procedure will call the 'regular' GVNS when going down through the different levels. This must be done to make sure that only solutions that are evaluated on each level are marked as promising and thus no simulation efforts are spent on non-fully evaluated solutions.

Algorithm 5: General Variable Neighbourhood Search with simulation

Function $GVNS(Sol, k_{max} = \{k_{max}^0, k_{max}^1 k_{max}^2\}, T_{max} = \{T_{max}^0, T_{max}^1, T_{max}^2\},$
$N = \{N_0, N_1, N_2\})$:
$operators = \{N_h, N_{sa}, N_{route}\} = \{N_0, N_1, N_2\}$
$Sol \leftarrow Intial\_Solution(Sol)$
$promisingSolutions \leftarrow \emptyset$
$T \leftarrow CPU\_Time(\,)$
$\boldsymbol{while}\ T \leq T_{max}^0\ \boldsymbol{do}$:
  $k \leftarrow 1$
  $\boldsymbol{while}\ k \leq k_{max}^0\ \boldsymbol{do}$:
    $k \leftarrow k + 1$
    $Sol' \leftarrow Shaking(Sol,\ k,\ \{N_h \cup N_{sa}\})$
    $Sol'' \leftarrow Multi\_Level\_VND(Sol',\ N)$
    $\boldsymbol{if}\ Sol''\ better\ than\ Sol\ \boldsymbol{then}$:
      $Sol \leftarrow Quick\_Simulation(Sol'')$
      $k \leftarrow 1$
      $promisingSolutions \leftarrow promisingSolutions \cup \{Sol\}$
  repeat
  $T \leftarrow CPU\_Time(\,)$
repeat
$bestSolutions \leftarrow Intesive\_Simulation(\{Sol_1, \ldots, Sol_m\} \subseteq sorted(promisingSolutions))$
$\boldsymbol{return}\ bestSolutions$

Both the quick and intensive simulations follow the same procedure, which is shown in Algorithm 6, which is a monte carlo simulation.

Algorithm 6: Simulation Procedure

Function $Simulation\ (Sol, num\_replications)$:
$performance \leftarrow \emptyset$
$\boldsymbol{for}\ rep = 1\ to\ num\_replications\ \boldsymbol{do}$:
  $Random.seed(\,) \leftarrow rep$
  $Sol' \leftarrow scenario(rep)$
  $Sol'' \leftarrow N_{route}(Sol')$
  $performance \leftarrow performance \cup \{Sol''\}$
end
$\boldsymbol{return}\ mean(performance)$

4.6 CONCLUSIONS ON MODELLING THE BEST NETWORK

In this chapter, we have formulated a mixed-integer linear program to mathematically describe how to find the best logistical network of Company X. To keep the model from growing to large, an aggregation of volumes and stops on a postal zone level has been applied, providing an upper bound on the number of stops equal to the number of postal zones in the Netherlands. However, even with the aggregation the number of variables that this model requires grows very large. Therefore, the model becomes intractable very quickly and is not useful to find the best possible network in the real-life case. Therefore, we have formulated a general variable neighbourhood search (GVNS) that connects decisions on the hub, service area, and routing level via a multi-level variable neighbourhood descent procedure.

Since the GVNS is a deterministic approach that, in principle, is not capable to capture the seasonal fluctuations in workload over time, we have expanded the GVNS to a simheuristic-based GVNS. In this simheuristic, promising solutions that are found while solving the deterministic GVNS are passed to a simulation procedure which tests the solutions' performance under multiple scenarios. The enables us to provide a final solution based on the performance over multiple scenarios.

5 NUMERICAL EXPERIMENTS

In this chapter, we want to find out what the best depot locations and accompanying service areas are for Company X to answer the fourth research question *'what is the best logistics network for Company X in the Netherlands?'.* In Section 5.1, the experimental design is discussed. In Section 5.2, we elaborate on the data instances used for different experiments. Section 5.3 will describe the results of the parameter tuning and model validation. In Section 5.4, the found best network for Company X is presented, while Sections 5.5 and 5.6 place these results into context by comparing them to other designs and methods (5.5) and by performing sensitivity analyses (5.6).

5.1 EXPERIMENTAL DESIGN

Before we discuss the experimental design, we first introduce the technical specifications of the hardware on which the experiments are performed. Experiments 4 and 5 are performed on a computer with an i5-1145G7 processor with a speed of 2.6 GHz and 16 GB of RAM. Experiments 1-3, 6, 7, 8, and 9 are performed on a virtual machine using sixteen Inter Xeon CPUs with a speed of 3.1 GHz and 64 GB of RAM. The GVNS algorithm is written in Python 3.10.8, using the Spyder 5.3.3 IDE.

To answer the fourth research question, '*what is the best logistics network for Company X in the Netherlands*?', and put the findings into context, several experiments need to be performed. Answering this question consists of four sub-questions, listed below, each with its own set of experiments of which the results should give the information required to answer them.

5.1.1 What are the best parameters for the algorithm, and can we validate the outcomes?

- **Experiment 1: Parameter Tuning**

In this experiment, we determine how many iterations for each level (k_{max}) are needed to achieve good convergence of the solution. Good convergence is considered to be when the objective value has stabilized. A trade-off has to be made between finding the best possible solution and the required computational time.

- **Experiment 2: Performance of different operators**

In this experiment, we test the performance of each operator. Performance is measured by comparing the objective of the GVNS without an operator with the performance of the GVNS with all operators while comparing the required computational time. Furthermore, we test in which order the operators should be used to get the best results.

- **Experiment 3: Impact of the initial solution**

In this experiment, we test different procedures to create an initial solution to find out which procedure is best to get high-quality solutions. The better the quality of the initial solution, the better the final solution the algorithm can find.

- **Experiment 4: Model Validation**

In this experiment, we compare the outcome of the routing used in the GVNS to the routing optimization software of company X to validate the accuracy of the used routing. For the goal of this research, it is important that the model correctly describes the effect of hub locations and service areas on the routing cost. Thus, for a set of depot & service area configurations, the routing costs from the model should show the same relative changes as the costs that are found by the optimization software.

- **Experiment 5: Required number of replications for quick & intensive simulation**

In this experiment, we determine how many replications are needed during the quick and intensive simulations to get a sufficiently small confidence interval. Since the quick simulations are done during the GVNS, the number of replications should be such that within a reasonable time, the smallest possible confidence interval can be computed. For the intensive simulation, the time required to perform the simulation is less of an issue because it is done after the GVNS has finished. We want to determine the number of replications required to get an as small as possible confidence interval.

5.1.2 What is the best network design for Company X using the simheuristic-based GVNS?

- **Experiment 6: Results of the simheuristic-based GVNS**

We test the simheuristic-based GVNS to determine the best depot locations and service areas. We use the 30 chosen possible hub locations and the median day scenario. The result is the network design with the lowest operational costs.

5.1.3 How does the found network design compare to other methods & designs?

- **Experiment 7: Performance compared to a simple heuristic**

In the simheuristic-based GVNS, the decision for hub locations and service areas is made by taking the routing costs into account. This enlarges the solution space for the algorithm which takes a lot more time to execute. Therefore, we want to know what the performance of the algorithm is when we use a simple allocation strategy for creating the service areas and thus only use the routing costs to determine the hub locations. This will require less computational power and thus allows for more time to explore different hub locations and might therefore be a good or better alternative to the simheuristic-based GVNS.

- **Experiment 8: Performance of initial network design**

In this experiment, we determine what the network costs would be if we take the current network design as input. The algorithm will then search for the best service area allocation for the given depot locations. This experiment provides a baseline for the network costs to compare the results of experiments 6 and 7 to.

5.1.4 What is the sensitivity of input parameters?

Experiment 9: Sensitivity analysis on the depot locations and other costs

To better understand the behaviour of the algorithm, we test to which extent the solution changes if the values of input parameters are changed. This enables us to get insights into the impact and sensitivity of certain parameters. We test the effects of changing the hub opening costs, fixed vehicle costs, variable cost per km, hourly costs for vehicles, time windows for different processes, and linehaul costs. Each parameter is tested for 50%, 75%, 100%, 125%, and 150% as compared to the values that are used in the earlier experiments.

5.2 DATA INSTANCES

For the experiments in this chapter, different data instances are used. Because of the large size of the problem, it is not feasible to test every experiment on the full problem. Therefore, instances of different sizes are created which are explained in the section.

5.2.1 Small test instances

The experiments that are performed for tuning the algorithm settings will be run using the test data instances. These instances consist of postal zones that are drawn randomly from the large historical dataset and thus can be located anywhere in the Netherlands. Because too small instances lead to solutions with only one opened hub, which is not representative and does not allow for good testing of the operators, we need to use instances that will need to open multiple hubs. We found that instances that consist of 300 postal zones are not too small to be representative but also do not become too large to be unpractical. In these instances, the ten hub locations that are in the initial plan are taken as possible locations. The volumes of each postal zone are taken out of a relatively busy day from the historical dataset.

These test data instances are used for experiments 1-3.

5.2.2 Large test instance

Additional to the small test data instances a larger instance is required. This instance is a set that consists of 1054 postal zones which make up the more densely populated part in the west of the Netherlands. In this set, multiple days as well as different combinations of four opened hub locations are used to create different service areas. These service areas are used in experiment 4 to validate the routing accuracy in the model and 5 in combination with the simulation instances (Section 5.2.4) to determine the required number of replications.

5.2.3 Real average instance

To determine the best network for Company X, we need an instance that describes the distribution of workload over the Netherlands as best as possible. For the deterministic part, the GVNS, we use an average scenario so that the GVNS does not underestimate nor overestimate the network costs. We have chosen the median day in terms of total volume over all postal zones from our historical dataset and used that day has the average scenario in the GVNS. The disadvantage of this approach is that not every postal zone is included and therefore, the allocation of service areas is biased towards this specific scenario. This bias is countered to some extent by the simulation procedures but is not fully erased.

Although, we consider each postal zone, taking the average workload for each logistical process for each postal zone is a very poor approach. This approach results in a situation in which every postal zone has to be visited because each has a nonzero average. This results in many routes that need to be driven and thus a large overestimation of the total network costs compared to an average day based only on the total volume of that day.

5.2.4 Simulation instances

A difficulty is that there exist great seasonal fluctuations between days, and thus finding a dataset that grasps these fluctuations well is (nearly) impossible. Therefore, to deal with these fluctuations, the GVNS is expanded with a simulation component to a simheuristic. For this simulation part of the simheuristic-based GVNS, it is necessary to create scenarios. Ideally, the scenarios are created based on probability distributions. However, the scenarios consist of multiple interrelated variables per postal zone, making fitting probability distributions a large task. In this research, we must scope and are not able to do the work required to do a complete simulation study.

Therefore, for this research, we use a large historical dataset from the e-commerce company out of which real days can be drawn that can be used as scenarios for the future. Since Company X is a start-up, it will need to build up the number of postal zones it services and thus also the total volumes that it transports. Therefore, it can be reasonably assumed that these scenarios form an accurate prediction for future workload because the total volumes in these scenarios will be reached by Company X in a few years. Moreover, these scenarios consist of real days that have happened in the past, making them accurate representations of reality.

The simheuristic expansion of the GVNS involves two parts: the quick simulation and the intensive simulation. The first is meant to create a quick assessment of the quality of promising solutions the GVNS comes across, whereas the intensive simulation is meant to give a thorough assessment of the best solutions found during the GVNS search. Therefore, we need a small number of scenarios for the quick simulation and a large number of scenarios for the intensive simulation.

The creation of scenarios is as follows, also shown in Algorithm 7:

1. Taking the historical dataset and determining the total volume of sales per day.
2. Filter out non-representative days that have less than 3.000 parcels as these are not expected to occur in the future anymore.
3. Sort the days by total volume.
4. Divide the days into buckets, the number of buckets is equal to the number of scenarios required for the quick or intensive simulation.
5. From each bucket, randomly choose a day to be used for a scenario.
6. For each day determine:
    a. The volumes per process per postal zone
    b. The number of stops per process per postal zone
    c. The volumes that need to travel between each pair of postal zones

Algorithm 7: Creation of simulation scenarios

Function $CreateSimulationInstances(historicalData)$:
$dataSet \leftarrow DetermineTotalVolumePerDay(historicalData)$
$dataSet \leftarrow Filter(dataSet, totalVolume, 3000)$
$dataSet \leftarrow Sort(dataSet, totalVolume)$
$dataSet \leftarrow CreateBuckets(dataSet, totalVolume, num\_scenarios)$
$scenarios \leftarrow \emptyset$
$\boldsymbol{for}\ b = 1\ \boldsymbol{to}\ num\_scenarios\ \boldsymbol{do}$:
    $day \leftarrow random.choice(dataSet, dataSet_{bucket} = b)$
    $singleScenario \leftarrow \emptyset$
    $\boldsymbol{for}\ z = 1\ \boldsymbol{to}\ |zones|\ \boldsymbol{do}$:
        $volume_z^p \leftarrow DetermineProcessVolume(day, z, p) \quad \forall\, p \in P$
        $numStops_z^p \leftarrow DetermineNumberOfStops(day, z, p) \quad \forall\, p \in P$
        $shipment_{z,d} \leftarrow DetermineShipmentVolumes(day, z, d) \quad \forall\, d \in zones$
        $singleScenario \leftarrow singleScenario \cup \{volume_z^{p \in P}, numStops_z^{p \in P}, shipment_{z,d \in zones}\}$
    end
    $scenarios \leftarrow scenarios \cup \{scen\}$
end
$\boldsymbol{return}\ scenarios$

5.3 MODEL TUNING, TESTING, AND VALIDATION

In this section, the experiments related to the tuning and performance of the model are discussed.

5.3.1 Experiment 1: Parameter tuning

Important parameters in the GVNS algorithm are the values of k_{max}, which determine the level of diversification that will be used to avoid the traps of local optima. The proposed GVNS framework operates on multiple levels, the hub and service area level, and thus a combination of two parameters must be chosen. It is known that for large values of k_{max} the VNS tends to behave like a random multi-start heuristic, i.e., a heuristic that randomly restarts without the features of the current best solution (Ratli et al., 2022). Yet, we still want to find high-quality solutions meaning the algorithm must be able to escape from local optima.

To test which combination of k_{max} values for both the hub level ($k_{\max}^0$) and service area level ($k_{\max}^1$) is the best, the GVNS is tested on the different small test scenarios (see Section 5.2.1). For each run, the same initial solution is chosen, so that the quality of the initial solution has no impact on the outcome of the GVNS. For each combination of k_{max} values, ten scenarios are solved. The results are summarised in Table 6.

Table 6 Results of experiment 1: Parameter tuning

			Objective				Time (s)			
Exp	$\boldsymbol{k_{\max}^0}$	$\boldsymbol{k_{\max}^1}$	**Min**	**Max**	**Mean**	**St. Dev.**	**Min**	**Max**	**Mean**	**St. Dev.**
1.1	3	3	6,341	8,591	7,479	708	36	138	73	35
1.2	3	5	6,748	8,531	7,581	638	43	252	125	64
1.3	3	8	6,942	8,314	7,541	526	102	330	204	75
1.4	5	3	6,412	8,603	7,450	647	57	351	147	94
1.5	5	5	6,624	8,509	7,556	610	113	468	229	106
1.6	5	8	6,038	8,134	7,387	615	67	440	273	116
1.7	8	3	6,423	8,096	7,375	605	105	254	173	48
1.8	8	5	6,255	8,337	7,459	663	146	461	276	102
1.9	8	8	6,387	8,366	7,555	568	206	627	393	155

Based on the performance on the objective value, it becomes clear that experiments 1.6 and 1.7 are the top performers on the mean objective value, with values below 7,400. Furthermore, looking at the run times the best-performing settings are those of experiments 1.1, 1.2, and 1.4, which all have a mean run time under 150 seconds. For the run time, this result is expected since these are the experiments that have the lowest levels of diversification. For the mean objective, the results show that the best solutions are not found using the highest levels of diversification, which seems to be in line with the statement of Ratli et al. to not have too large values for k_{max}. Concluding, we will use the settings of experiment 7 for the remainder of this research. These settings yield the best average result, and they do provide good results in a reasonable time.

5.3.2 Experiment 2: Operator performance

To understand the functioning of the model, we investigate the performance of every single operator. To test this, we have defined nine experiments in which a single operator is excluded. The first experiment includes all operators to provide a benchmark to compare the other experiments, note that this is experiment 1.7.

Each experiment is performed ten times, using the small test data instances (Section 5.2.1). For each experiment the minimum, maximum, mean, and standard deviation of both the objective value and run time are determined. The results of the experiments are presented in Table 7.

Table 7 Results of experiment 2: Operator performance

	Missing	Objective				Time (s)			
Exp	**Operator**	**Min**	**Max**	**Mean**	**St. Dev.**	**Min**	**Max**	**Mean**	**St. Dev**
1.7	None	6,423	8,096	7,375	605	105	254	173	48
2.1	add_hub	6,830	8,607	7,580	661	67	171	114	35
2.2	remove_hub	7,062	8,899	8,340	581	11	30	18	5
2.3	swap_hubs	6,398	8,537	7,634	636	37	210	95	52
2.4	move_zones_PC2	6,694	8,053	7,432	447	67	235	118	55
2.5	move_zones_PC3	6,388	8,049	7,472	518	54	214	116	59
2.6	swap_zones	6,342	8,784	7,690	718	53	311	127	77
2.7	move_zone_Allp	6,231	8,468	7,532	711	60	210	129	54
2.8	move_zone_1p	6,748	8,186	7,487	496	47	172	104	38

There are a few things that stand out from these results. The mean objective of experiment 2.2 is much higher than those of the other experiments. Furthermore, the operator move_zones_PC2 has the lowest objective value, followed by the move_zones_PC3 operator.

The high average value for the remove_hub operators is understandable. The test is performed on a small instance of 300 postal zones, and from the results, we see that most final solutions have only three hubs opened, while the initial solution has six. Therefore, not being able to remove a hub disables the algorithm from finding better solutions, hence also the low average run times. Little improvements mean that the algorithm progresses quicker since there is no reset in the diversification parameter k.

Less expected is the lower value for the move_zones_PC2 and move_zones_PC3 operators. Understandably, the impact of all operators on the service area level (experiments 2.4-2.8) is lower than those of the hub-level operators since hub-level operators have a bigger impact on the network. However, between the service area operators, we would expect the move_zones_PC2 and move_zones_PC3 operators to have a bigger impact on the objective than the operators that move single postal zones, because they move larger areas at once. The result implies that the improvement in service areas lies mostly in the replacements of single postal zones and less in the movement of multiple.

The last important observation is that the average performance of all experiments is worse than that of the experiment with no excluded operators. This means that from these experiments no operator has shown such performance that we should exclude it from the algorithm.

Another element in evaluating the performance of the operators is by assessing the order in which they are applied. Since there can be variations on two levels with three and five operators respectively, the total possible combinations that would need to be assessed is $3! + 5! = 6 + 120 = 126$. This number is too large to assess all the possibilities within the timeframe of this research. We can however test the hub level as this concerns only 6 combinations.

Different orders for the 'add_hub', 'remove_hub', and 'swap_hubs' operators (AH, RH, SH resp.) are tested while keeping the order of the lower-level operators equal. Each experiment is tested for the ten small test instances (see Section 5.2.1), using no time restrictions, a fixed initial solution, and k_{max} settings as found in experiment 1. The results are presented in Table 8.

Table 8 Results of experiment 2: Operator order

Exp	Operator Order	Objective Min	Max	Mean	St. Dev.	Time (s) Min	Max	Mean	St. Dev.
2.9	AH-RH-SH	6,051	8,091	7,292	614	88	337	188	80
2.10	AH-SH-RH	6,749	8,804	7,583	677	77	288	159	66
2.11	RH-AH-SH	6,399	8,334	7,452	613	50	217	110	57
2.12	RH-SH-AH	6,052	8,294	7,509	672	78	199	108	36
2.13	SW-AH-RH	6,556	8,101	7,488	543	59	251	129	55
2.14	SW-RH-AH	6,231	8,287	7,467	609	57	154	107	34

In Table 8, we can see that the performance differences are not very large. However, experiment 2.10 is the best-performing setup. It has the best mean objective value and has achieved the best-observed value for a single run. The run time of this setup is longer, but that is a consequence of finding a better solution, as that resets the diversification parameters.

Concluding, in this experiment we have seen some unexpected results. The impact operators move_zones_PC2 and move_zones_PC3 showed to be the lowest, while it was expected that these would have a larger impact. Regardless, the performance of each experiment in which an operator was excluded was worse than the performance of the experiment with all operators, meaning no operator can be excluded. Furthermore, the order in which the hub level operators are applied was tested and the order 'add_hub'-'remove_hub'-'swap_hubs' was found to be the best performing one. Therefore, in further experiments, the operators (and their order) will be applied as introduced in Section 4.4.2.

5.3.3 Experiment 3: Impact initial solution

The quality of the initial solution can have a great impact on the final solution of the GVNS. Having a bad starting solution means that the algorithm has to spend time getting to a reasonable solution while having a good starting solution allows the algorithm to spend that same time on finding a really good solution. Therefore, we will investigate the impact of using different procedures to determine an initial solution. As explained in Section 4.4.1.1, we have three different procedures: the random solution, the greedy clustering and the improved clustering.

In this experiment, we test each procedure for the ten small test instances (see Section 5.2.1). The values for k_{max} are [8, 3] as determined in experiment 1 and the run time is not limited. For both the clustering procedures a value for the number of hubs k needs to be given as input, therefore we test these procedures with the values of 4, 5, 6, 7, and 8 for parameter k. These are chosen by taking the ceiling of the average between the minimum number of hubs to open (1) and the maximum number of hubs (10), which is 6, and take a range of ± 2. The results of these runs are shown in Table 9

Table 9 Results of experiment 3: Impact of the initial solution

		Objective				Time (s)			
Exp	Initial Solution	Min	Max	Mean	St. Dev	Min	Max	Mean	St. Dev.
3.1	Random	6,535	8,704	7,532	695	68	322	153	76
3.2	GrdCls_k4	6,514	8,029	7,367	523	74	251	139	62
3.3	GrdCls_k5	6,586	8,133	7,507	530	86	242	147	49
3.4	GrdCls_k6	6,760	8,271	7,467	547	82	242	147	50
3.5	GrdCls_k7	6,231	8,401	7,593	642	62	240	138	64
3.6	GrdCls_k8	6,608	8,274	7,553	520	64	389	175	91
3.7	ImprCls_k4	6,342	8,065	7,375	586	84	203	126	39
3.8	ImprCls_k5	6,342	8,291	7,532	596	44	362	168	95
3.9	ImprCls_k6	6,231	8,268	7,382	614	49	313	151	75
3.10	ImprCls_k7	6,301	8,159	7,440	630	115	300	196	54
3.11	ImprCls_k8	6,342	8,545	7,552	639	109	233	170	35
	Avg. GrdCls	6540	8222	7497	552	74	273	149	63
	Avg. ImprCls	6312	8265	7456	613	80	283	162	60

From Table 9 it becomes clear that the random initial solution does not have the best mean objective performance and has the biggest variance (as can be expected). Therefore, this procedure for creating an initial solution is not deemed to be a good one. The greedy clustering procedure has mixed performance, both the best and worst mean objectives are found using the greedy clustering procedure. Also, the four best performances in terms of standard deviation are done by the greedy clustering. For the improved clustering procedure, we can see that the second, third, and fourth performances on the mean objective are the result of the improved clustering. Looking at the averages over the five greedy clustering experiments and the five improved clustering experiments, we can see that the improved clustering approach has a better average mean objective and also has a better average minimum mean objective.

However, the detailed output (see Appendix C) showed that the initial solution generated by the greedy and improved clustering procedure is the same for each scenario. Thus, the better performance of the improved clustering cannot be explained by the starting solution since these are the same. It turns out that for seven opened hubs (exp. 3.10) the runs of the improved clustering procedure were able to reach better solutions, which largely explains the slightly better average performance of the improved clustering. Since the best solutions have either three or four opened hubs, the algorithm has to spend more time on improving the bad start solution of seven opened hubs and thus has a lower probability of finding good solutions. Due to the randomness in the search, the algorithm happened to find some better solutions when the initial solution was made using improved clustering.

Concluding, although the performance of both clustering procedures is quite similar and the starting solutions are similar for each scenario, we will use the improved clustering. For larger instances, the improved clustering procedure could improve the initial solution compared to the greedy procedure.

5.3.4 Experiment 4: Model validation

Company X makes use of a route optimization software package to plan the daily routes. By using this software package, we can create estimates of future route lengths, durations, and the number of routes. These estimates can be used to test how accurate the routing costs are that the GVNS model takes into account when calculating the network costs.

For this experiment, we use the larger test instances consisting of 1054 postal zones and considered four possible hub locations (see 5.2.2). Based on 10 days from the historical sales dataset, we created different scenarios in which one or more hubs are opened and each postal zone is serviced from an opened hub. The route optimization software gets a volume to be transported per postal zone and a set of opened hubs. The software determines which postal zone to service from which hub while searching for the minimal costs of routes. The output of the software is a list with routes, from this list we can determine for each opened hub: which zones the hub services; how many routes are used; what the total travelled distance is; and what the total duration of the routes together is. This resulted in 165 hub-service area combinations and corresponding route distances, durations and number of routes.

Consecutively, we used the Clarke & Wright savings algorithm to determine routes using the same hub-service area combinations. This resulted in a total route distance, total route duration, and a number of routes to compare with the outcomes of the route optimization software. For these three characteristics the mean percentage error (MPE) and mean absolution percentage error (MAPE) are determined. Since finding the best possible routing is not the goal of this model, we also determine the MPE and MAPE on the estimates relative to the mean of both the routing software as the Clarke & Wright heuristic. The idea is that when creating service areas using routing costs, the exact costs are as not as relevant as accurately reflecting changes in the service area in the routing costs. For example, if the routing software estimates the total distance at 1.5 times the average routing distance, the heuristic should also result in a distance of approximately 1.5 the average distance. The MPE and MAPE are calculated by equations (46) and (47), respectively, where S_i indicates the value for distance, duration, or number of routes as determined by the optimization software, and H_i the values found by the heuristic for the corresponding instance.

$$MPE = \frac{1}{k}\sum_{i=1}^{k}\frac{S_i - H_i}{H_i} \cdot 100\% \qquad (46)$$

$$MAPE = \frac{1}{k}\sum_{i=1}^{k}\frac{|S_i - H_i|}{S_i} \cdot 100\% \qquad (47)$$

The results are summarized in Table 10. First of all, there is some difference between the results of the optimization software and the heuristic. This is expected since Clarke & Wright is a greedy constructive heuristic and thus would not perform as well as the optimization software package. The negative MPE values show that the distance and duration found by the heuristic are generally higher than those of the optimization software. However, the deviations for the duration of the routes are not very high. Because the duration of routes is the major cost component, this implies that the heuristic method is suitable for our application. When looking at the performance relative to the mean better performance can be seen for both the distance and duration of routes. This shows that the heuristic does respond quite similarly to changes in service areas as the optimization software. Furthermore, based on the outcomes of both methods the corresponding costs are calculated. Especially the MPE for costs to mean is very low, implying that on average the outcomes of both methods in terms of costs change similarly when service areas change.

Table 10 Results of experiment 4: Model validation

Performance C&W v. routing software	**MPE**	**MAPE**
Distance	-0.4514	0.4956
Duration	-0.0506	0.0988
Number of Routes	0.1099	0.4836
Costs	-0.0842	0.2272
Distance to Mean	-0.0349	0.2287
Duration to Mean	0.05	0.086
Num. Routes to Mean	0.1721	0.4878
Costs To Mean	-1.884e-17	0.2082

All in all, because of the high values for the MPE and MAPE, it can be concluded that the Clarke & Wright savings heuristic does not perform as well as the optimization software, but this is expected. However, for the performance relative to the mean, the differences are much smaller. Therefore, it does reflect the changing service areas well enough to be used to assess the impact of service areas on the routing costs and can be used in our GVNS model.

5.3.5 Experiment 5: Required number of simulations

During the quick and intensive simulation modules of the simheuristic, the mean and average performance of a solution are determined. For these results to be of high enough quality, we must find the minimum required number of replications that need to be performed to get results with high enough significance. The quick simulation is performed during the GVNS, therefore we want to get a reasonable level of significance (e.g. 90%

or 95%) while not having to spend too much computation time. For the intensive simulation, we want to have a higher level of significance and the required computation time can be longer because these simulations are performed after the GVNS procedure. However, because the historical dataset includes about 500 days, there is a limit on the number of replications we can run. Therefore, we have to find the highest level of significance under 500 replications.

To determine the required number of replications, the minimum number of replications must be found for which the half width of the confidence interval (δ) relative to the mean ($\bar{X}$) is smaller than the relative error γ' (Law, 2014). The mean and confidence interval are found by taking a given solution, i.e., fixed hub and service area allocation, for which 500 replications are performed. In each replication, a random day from the historical dataset is chosen and the routing is calculated using the Clarke and Wright heuristic after which the total network costs can be calculated.

For each replication i, the mean network costs and the sample variance are calculated by taking $\bar{X} = \frac{1}{i} \cdot \sum_{k=1}^{k=i} x_k$ and $S^2 = \frac{1}{i} \cdot \sum_{k=1}^{k=i} (x_k - \bar{X})^2$ respectively. Then for each replication, it is checked whether equation (48) holds. In this equation, the relative error is given by $\gamma' = \frac{\gamma}{1+\gamma}$ and the confidence interval half-width is given by $\delta = t_{n-1,1-\frac{\alpha}{2}} \cdot \sqrt{\frac{S^2}{n}}$. The replication number for which the equation holds for both itself and all following replications is the minimum number of required replications.

$$\frac{t_{n-1,1-\frac{\alpha}{2}} \cdot \sqrt{\frac{S^2}{n}}}{\bar{X}} < \gamma' \qquad (48)$$

We have tested this equation for three confidence levels α = 0.05, α = 0.02, and α = 0.01. For each level of alpha holds that $\gamma = \alpha$. The results are shown in Figure 23 Confidence interval half-width relative to the mean for α = 0.05, Figure 24 Confidence interval half-width relative to the mean for α = 0.02, and Figure 25 Confidence interval half-width relative to the mean for α = 0.01. From the figures, we can conclude that after 12 replications we can achieve a relative error on the average network costs of 0.05 at a significance level of 95%. Similarly, we find that for a significance level of 98%, we achieve a relative error of 0.02 after 106 replications and that after 500 replications the relative error has not passed below the required level and thus a significance of 99% cannot be achieved.

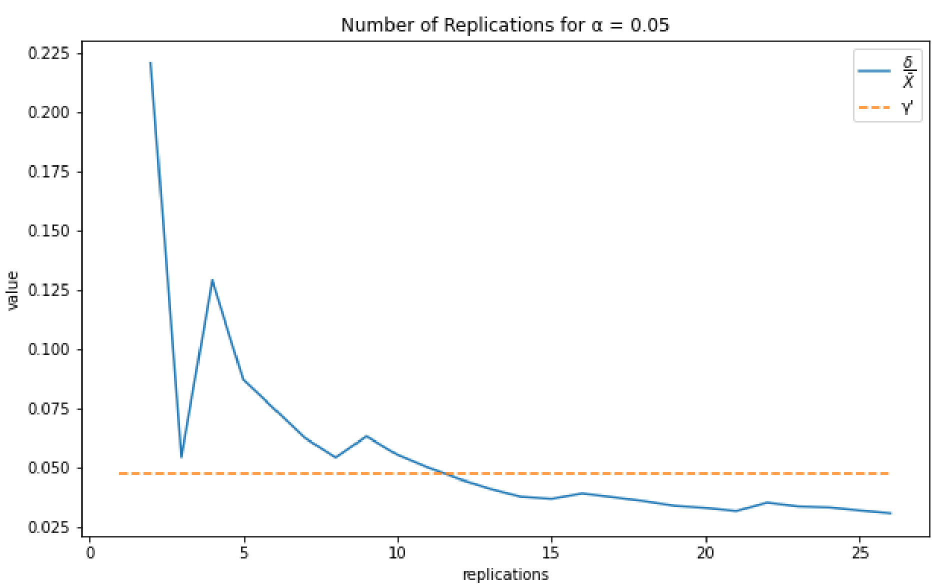

Figure 23 Confidence interval half-width relative to the mean for α = 0.05

Performing 12 replications takes around 30 seconds, which is relatively short. Meaning that this is a suitable number of replications for the quick simulation that is performed during the GVNS. Furthermore, we can achieve a significance of 98% by performing 106 replications. This takes little under 300 seconds and is, therefore, suitable for intensive simulation that is performed when the GVNS is finished.

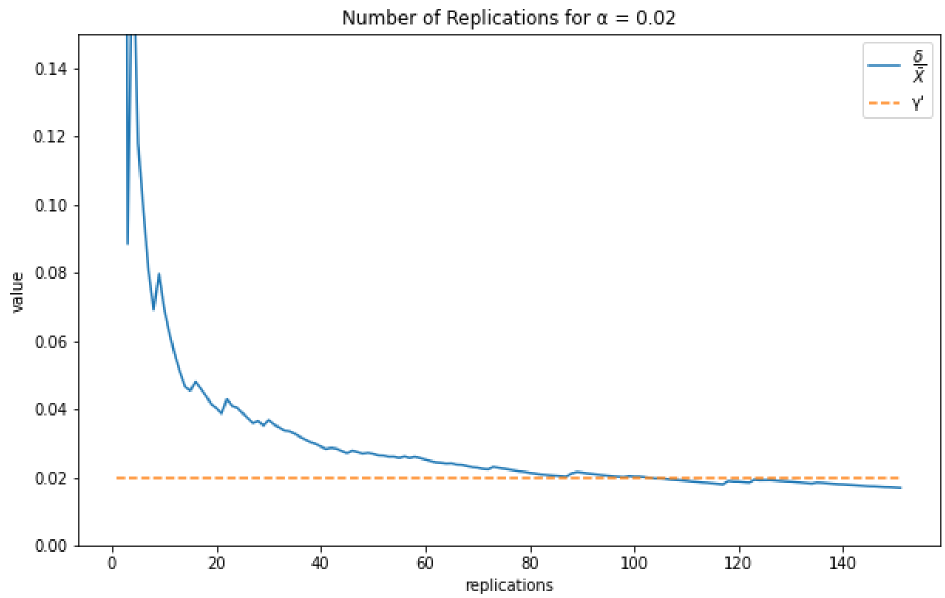

Figure 24 Confidence interval half-width relative to the mean for α = 0.02

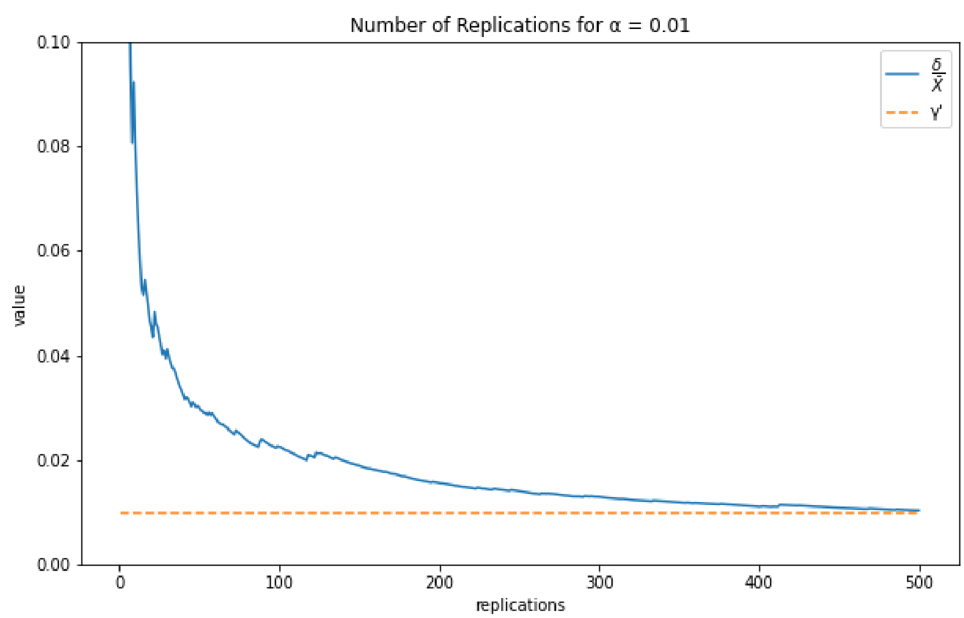

Figure 25 Confidence interval half-width relative to the mean for α = 0.01

5.4 RESULTS OF SIMHEURISTIC-BASED GVNS

In this experiment, we evaluate the results of running the proposed simheuristic-based GVNS on the real average instance to find the best network for Company X. The model is run with the settings as found in experiment 1, thus we use the values 8 and 3 for $k^0_{\max}$ and $k^1_{\max}$ respectively. The initial solution is created using 10 opened hubs, similar to the number in the initial design. Furthermore, because of time restrictions for doing this research, the model run time is limited to 48 hours to have results on time. Because of the strategic nature of the problem, the model does not have to be used very frequently and thus letting it run for a long time is no problem in practice.

Figure 26 Network design determined by the simheuristic-based GVNS.

During the GVNS procedure, no promising solutions were found, meaning that the resulting solution is the initial solution. This solution was investigated more thoroughly during the intensive simulation. The resulting total network cost is on average €25,290 and has a standard deviation of €5,244.

Looking at the solution, see Figure 26 Network design determined by the simheuristic-based GVNS, we see the ten opened hub locations. Three of these hubs are in locations we would not expect them to be. Firstly, in Section 2.2.3.2, we found that in the northern parts of the Netherlands, there is a relatively low demand per postal zone. However, three hubs are opened in the north, which we expect to be more than required. Secondly, the hub in the purple area is located quite northerly while we would expect it to be located near Amsterdam since there the workload is much higher. Lastly, the hub in the Zeeland province (lower left). This is an area with low population density and thus less workload. We would expect that moving this hub more westwards would yield a better solution.

To test whether these changes would improve the solution, we manually removed the upper left hub, moved the hub with the purple service area near Amsterdam, and the Zeeland hub more westwards. This solution was investigated by putting it into the intensive simulation, which resulted in an average network cost of €24,739 and a standard deviation of €5,387. Thus, in contrast to our expectations, the changed solution performs worse.

Concluding, the simheuristic-based GVNS has not shown to be capable of finding better solutions than its initial solution. This initial solution has an average network cost of €25,290. Based on the workload analysis in Section 2.2.3.4, we expected that we could manually improve this solution. However, based on the simulation results, we were not able to do so.

5.5 RESULTS COMPARED WITH OTHER METHODS

To be able to assess the added value of our proposed simheuristic-based GVNS approach, we need to compare its results to other methods. Therefore, we tested what the resulting network design would be if we would use a simple service area allocation method and we tested what the network costs would be with the hub locations of the initial plan. The difference in objective between these results and the simheuristic-based GVNS indicates its added value for Company X.

5.5.1 Experiment 7: Comparison of GVNS to a simple heuristic

In this experiment, we test how the simheuristic-based GVNS compares to a method with a simple service area allocation method. Since this simple allocation is much less expensive in terms of computational efforts, it is useful for Company X to know whether similar or better results can be attained using the simple approach.

We use the same GVNS framework, however, we do not use the operators on the service area level as presented in Section 4.4.2, but replace those with a single operator that will assign each postal zone to the opened hub that is closest by. This means that we implicitly assume that there exists a direct connection between each postal zone and the hub it is allocated to, making this approach similar to solving an HLP. However, in contrast to an HLP, the routing costs are determined still to make sure we can compare its outcomes to those of the simheuristic-based GVNS.

The model is run with the same settings as the simheuristic-based GVNS. During the GVNS, one promising solution was found which was put into an intensive simulation. The resulting total network cost is on average €24,739 and with a standard deviation of €5,259. This network design is shown Figure 27. In this experiment, the initial solution, which is the result of experiment 6, was improved by removing a hub in the north.

Figure 27 Network design using the simple allocation method.

We can conclude that the simple approach, under these settings, performs slightly better than our proposed simheuristic-based approach. Due to the simple allocation strategy for postal zones, less time is spent on finding better service areas and thus more time can be spent by the algorithm to adjust the hub locations. However, only closing one hub and not finding other promising solutions is not much of an improvement compared to the initial solution. Also, in experiment 6 we did not see an improvement. This could indicate that the search process on the hub level is not effective enough at the moment.

Therefore, we have performed some additional tests where we have tested the simple approach for the ten test instances to be able to compare the simheuristic-based GVNS and the simple approach with each other, see Table 11. From these results, we can conclude that the performance of both approaches is very similar for the test instances with the simheuristic-based GVNS having a better mean objective.

Table 11 Results of the simple approach compared to simheuristic-based GVNS for test instances

		Objective				Time (s)			
Exp	**Approach**	**Min**	**Max**	**Mean**	**St. Dev.**	**Min**	**Max**	**Mean**	**St. Dev**
1.7	SH-GVNS	6,423	8,096	7,375	605	105	254	173	48
7.2	Simple Method	6,236	8,238	7,377	645	51	146	91	27

Concluding, using the simple allocation strategy we have found a solution that is slightly better than that of the simheuristic-based GVNS. This solution is found by closing one hub compared to the initial solution. Because both the GVNS and the simple approach do not find other promising solutions, we expect that the search on the hub level is not effective enough. To be able to compare the performances of both methods we applied the simple approach to the test instances and compared those results with the outcomes of experiment 1.7. From these results, we can conclude that the simheuristic-based GVNS performs better.

5.5.2 Experiment 8: Comparison to initial network design

To determine the quality and added value of our simheuristic-based GVNS approach, we determined the performance of the initial network design. In this initial network, ten hub locations have been determined based on demand-weighted clustering rules (see Section 2.4).

For this experiment, we have used simheuristic-based GVNS but left out the hub level. The ten hub locations of the initial design are given as input to the model, after which the service area and routing operators are applied to find the best possible service area allocation for these ten hubs.

Because the hub level is excluded, the settings for the values of k_{max} as determined in experiment 1 are not applicable. Instead of using a k_{max} of 3 we chose to set the value of k_{max} to 5, to allow for some more diversification so that there is a greater likelihood of finding the best service areas for the hub locations.

The GVNS procedure found promising solutions, which were investigated more thoroughly in the intensive simulation. The six solutions performed very similarly, with the average over all scenarios ranging from €27,261 to €27,297[7]. Furthermore, we found that the standard deviation increased with the average value of each solution, so there is one solution which performs best in both the average as well as the variance over the network costs.

The objective that is found while improving the initial network design is significantly higher than that of the simheuristic-based GVNS (+€1,972) and the simple allocation strategy (+€2,522). Meaning both those methods can find a better network design than the initial plan.

5.6 SENSITIVITY ANALYSIS

In this experiment, we test the impact of different inputs on the objective value of the GVNS. By doing so we get a better understanding of the impact of certain parameters and can test how sensitive the model is to these input parameters. If the model is sensitive to changes in certain input parameters, we know that Company X must either avoid or attend to changing that parameter.

We have looked at different input parameters related to hubs, vehicles, process characteristics, and linehaul: hub opening cost; PPC duration; HD duration; fixed vehicle cost; hourly vehicle cost; cost per km; vehicle capacity; vehicle driving range; and the linehaul cost factor. Each parameter is tested for five levels 50%, 75%, 100%, 125%, and 150%, each level is relative to the values that are used in the previous experiments. While changing the value of one parameter, all the other parameters are kept equal. Therefore, the settings for the 100% case are the same for every single experiment.

Furthermore, these experiments are run using the test instances of 300 postal zones and 10 possible hub locations (see Section 5.2.1). To better test the impact of the input parameters we fixed the starting solution for each run, in this way the quality of the starting solution cannot influence the performance of each run. Due to time constraints for performing this research, each experiment setting is only performed once. The impact of changing the different input parameters on the objective value can be found in Table 12. Apart from the impact on the objective value, we looked at the effects on the number of opened hubs, total route length, total route duration, the total number of routes, and linehaul costs. These results can be found in Appendix F.

[7] *The solutions found in this experiment can be seen in Appendix E.*

Table 12 Results of experiment 9: Sensitivity Analysis

		Objective				
Exp	**Input parameter**	**50%**	**75%**	**100%**	**125%**	**150%**
9.1	Hub opening cost	6,692	7,255	7,861	8,444	8,941
9.2	PPC duration	8,139	8,060	7,956	7,513	8,248
9.3	HD duration	8,424	7,854	8,189	7,579	7,643
9.4	Fixed vehicle cost	7,344	8,129	7,956	7,786	7,937
9.5	Hourly vehicle cost	6,458	7,158	7,689	8,283	9,040
9.6	Cost per km	7,678	7,796	7,825	7,866	7,691
9.7	Vehicle capacity	7,876	8,228	7,634	7,894	7,989
9.8	Vehicle driving range	8,132	8,188	7,642	7,560	7,711
9.9	Linehaul cost factor	7,685	7,688	7,956	7,602	7,964

The resulting objective values, which can be seen in Table 12, are not changing monotonically to the change of the investigated input factor. There are two possible causes that should be considered when looking at these results. Firstly, the experiments are only performed once and thus it can happen that due to randomness the presented outcome is a bad one. Secondly, the results in Table 12 should be looked at in combination with the results presented in Appendix F. Because of the large number of variables, the increase in cost from one parameter can be 'compensated' by making different choices on other cost factors.

Furthermore, from both Table 12 and Appendix F it becomes clear that the two largest cost factors are the hub opening costs and routing costs, the linehaul costs are almost neglectable. It must be noted that in these objective values, the penalty costs are included. These are relatively high because the zones are spread over the Netherlands but with only 300 there is no reason to open more hubs and thus there will be large distance allocations. Looking at individual parameters, we see that there are four that have a big impact on the total costs. These are the hub opening cost, fixed vehicle cost, vehicle cost per hour, and cost per km. Especially the cost per hour has a very big impact on the routing costs that range from 2,258 to 5,218 for the costs ranging from 50% to 150%. This effect clearly shows in the objective, the range of the objective value is the largest for the cost per hour. The parameter that has the second biggest impact is the hub opening costs. When the hubs are cheaper, more are opened which leads to lower routing costs.

While the changes in the fixed, km, and hourly costs for vehicles are represented in an expected way in the routing costs, the total route length, duration, and the number of routes do not respond similarly. Even more so, for the highest values of cost per km and cost per hour, the largest route lengths and durations are observed. We believe this is caused by the Clarke & Wright heuristic. In building the routes, the heuristic does not take the cost parameters into account. These are used only to create the cost and savings matrix, but these will not change the order in which stops are added to the routes since all values in these matrices change proportionally to the parameter value. The heuristic does take capacity, time, and driving range constraints into account. For changes in the time window (duration PPC, HD) and driving range parameters, we can see a monotonic change in route lengths, duration, and number of routes. For the capacity parameter, there is no clear pattern visible in the resulting route length, duration, and number of routes. We expect that for the used test instance, the vehicle capacity is not a constraining factor and therefore does not show the same results as time windows and driving range.

Overall, we can see that the network costs consist mainly of routing and hub opening costs. Furthermore, we have seen that some parameters have a large impact on the total network cost, the biggest are the cost per hour for a vehicle and the costs of opening a hub. When Company X can keep these as low as possible, they can lower their network costs. The constraining factors of time windows and driving range also have a negative impact on the routing costs, so the larger they are the lower the routing costs. We would also expect this behaviour from the vehicle capacity, however, from these experiments we could not observe a clear relationship between the capacity and the routing costs, which might be due to the instance size.

5.7 CONCLUSION ON NUMERICAL EXPERIMENTS

In this chapter, we aimed to answer the question *‘what is the best logistics network for Company X in the Netherlands?’.* To do so, we first determined the model settings. We found that the settings that resulted in the best objective value were to use a k_{max} level of 8 and 3 for the hub and service area levels respectively. We also tested the performance of the operators and assessed the order in which they are applied on the hub level. We found that there is a difference in performance between the operators but using all of them yields better solutions than excluding one. On the hub level, the best order to apply the operators has been found to be ‘add_hub’-‘remove_hub’-’swap_hubs’ Furthermore, we found a discrepancy between the routing costs of our routing heuristic and the routing software of Company X. However, it does seem reasonably able to capture the effect of changing service areas on routing, making it useful enough for our model. Also, we determined that the required number of replications in the quick and intensive simulations needed to be 12 and 106 to attain confidence intervals of 95% and 98% respectively.

Looking at the network design resulting from the simheuristic-based GVNS, we see that the best network operates at €25,290 per day. Furthermore, we found that using a simple allocation method for the service areas performed slightly better at €24,739 and that the initial plan of Company X performs worse with on average €27,261 per day. Thus, the best network design has been found by the simple allocation method.

Furthermore, while performing and analysing both experiments 6 and 7, we found some improvement points to increase the efficiency of the model. Currently, not a lot of new, better-performing, solutions are found, which can be caused by the operators being not efficient enough. Therefore we recommend improving the operators to not only make changes randomly. To be able to compare the simheuristic-based GVNS and the simple allocation method, we compared them based on the ten small test instances and found that the simheuristic-based GVNS performed better.

From performing all the different experiments, we can draw another conclusion. The GVNS takes a deterministic scenario as input. We found that the choice for this scenario has a great impact. We would want to include each postal zone in this scenario since each could be serviced in one of the processes. However, doing so leads to a scenario that contains many (small) stops, leading to way higher routing costs than what can be seen when solving a scenario taken from the historical dataset. Therefore, we have chosen to use the median day, based on the total volume, as the deterministic scenario for the GVNS.

6 CONCLUSIONS & RECOMMENDATIONS

In this chapter, we summarise the research findings, draw conclusions and give recommendations for further research. We do so by answering the last research question *'What are the conclusions of the research and recommendations for Company X?'*.

6.1 CONCLUSIONS

In this research, we investigated how to model the most efficient and robust logistical network for Company X to enable them to achieve their ambition to be the most reliable, efficient, and sustainable parcel carrier in the Netherlands. To do so, the main research question is *'What is the best location of and allocation to depots such that the logistical network of Company X operates at minimal costs?*'. This model must be able to determine the network that minimizes the operational costs by including the costs for opening hubs, linehaul, and the cost of routes while also assessing its performance given multiple scenarios.

We found that in the literature the hub location routing problem is the framework that describes all these elements and is therefore suitable for this model. We modelled the network of Company X using a MILP formulation. Due to the size of the problem, using exact optimization becomes intractable. Therefore, we have proposed a heuristic approach that combines elements from the fastest method found in literature and the most suitable method for the context of Company X. The resulting heuristic is a general variable neighbourhood search that aims to find the best configuration on the hub level as well as the service area level. Using a deterministic GVNS does not capture the stochasticity that exists in the workload of Company X. Because of the strategic nature of the problem as well as the desire of Company X to have a robust network design, incorporating stochasticity is required. Therefore, we have expanded the GVNS approach into a simheuristic-based GVNS in which the most promising solutions found by the GVNS are simulated to assess their performance under multiple scenarios of parcel volumes.

In the numerical experiments, the simheuristic-based GVNS found a network that consists of ten hub locations and operates, on average, at €25,290 per day. The simple service area allocation method resulted in a network consisting of ten hub locations that operates, on average, at €24,739 per day, which is slightly better. The results were also compared to the initial plan for the network design that was made for Company X. This initial plan consists of ten hub locations and achieves an average performance of €27,261 per day, which is 8.4% worse than the simple method. Therefore, we can say that we have created a model that improves the operational cost of the network.

6.2 RECOMMENDATIONS FOR FUTURE RESEARCH AND LIMITATIONS

During this research, we found interesting results that will help Company X in the design of the logistics network. However, we also faced limitations and had to make decisions that could be solved in future works.

6.2.1 Limitations

A limitation of the chosen implementation of the GVNS framework lies in the different levels, namely the hub, service area and routing levels. We have chosen to implement separate GVNS on different levels of the problem. This allowed us to determine the best possible service area allocation for a given set of hub locations. Furthermore, using different levels makes the algorithm applicable to both strategic decisions (when hub locations are to be determined) and tactical decisions (when hub locations are fixed). The downside however is that the algorithm is spending computation time on improving the service area allocation of a bad set of hub locations. It would be more beneficial to, somehow, be able to assess the solution quality in an earlier stage so that computation time is spent more usefully.

Furthermore, assumptions are made for both the model and the input data. These assumptions have an impact on the outcomes of the model and thus can be a limitation. An example is the set of 30 possible hub locations that were generated. These determine the possible outcomes and choosing 30 other locations might have given us different results.

6.2.2 Recommendations for future work

Partly due to the above-mentioned limitations, the model and the tool built into this research can be further improved to provide more value to its users and insights into hub location routing problems. We have two recommendations that relate to the usage of the model and five recommendations to improve the model, which are listed in order of complexity.

In terms of using the model, the first recommendation is to research which deterministic scenario is best to use in the GVNS. We found that using a suitable scenario is of great importance since it can have a large impact on the network costs as well as impose a bias on the resulting solutions. Furthermore, a single scenario is never a perfect representation, therefore we recommend performing the GVNS on multiple different deterministic scenarios. This will result in a more elaborate list of promising solutions, which differ more strongly from each other, and thus a higher likelihood of finding the best logistical network. Secondly, in the experiments, only one improved solution was found. We recommend changing the termination criteria of the model such that while the time limit has not been reached, the diversification parameter k is reset so that the search continues.

With regard to improving the model, we have several recommendations. The first is to adjust the hub level operators to not only make changes randomly but also include some logic on which the choice for hubs to add, remove, or swap is based. By doing so, better solutions might be found before the termination criteria are met. The second recommendation is to model the hub opening cost in a more advanced way. In this research, we assumed a fixed opening cost for each hub independent of hub size. In reality, the cost of opening a hub is very much dependent on the size of the hub. Both in terms of asset costs as well as the cost of handling the volumes. Therefore, modelling the hub costs more accurately by, for example, piecewise linear formulations results in a more accurate model. The third recommendation is to draw the scenarios in the simulation procedure from probability distributions instead of using historical days as scenarios. This will make the model more scalable for the future and allows for flexibility in the scenarios that can be tested. The fourth recommendation is expanding the problem formulation as a multi-allocation problem. The current model formulation allows postal zones to be serviced from only one hub. However, it is interesting to investigate whether allowing postal zones to be serviced from multiple hubs leads to higher operational/routing efficiency due to better usage of vehicle capacity. The fifth recommendation is to find a better alternative for the Clarke & Wright heuristic to determine the routing cost. The major disadvantage of this method is that it requires more computation time when the problem size increases. On top of that, it is a constructive heuristic and thus will always yield an overestimation of the actual routing costs. Therefore, the model would perform better when a more accurate and scalable method would be implemented to determine the routing in this model. The last recommendation is that we recommend Company X to look into a way to incorporate the placing of and allocation to collection points in the routing. When the model is able to do so, it becomes possible to determine the best locations for these collection points which would further improve the efficiency of Company X's network.

6.3 CONTRIBUTIONS TO LITERATURE AND PRACTICE

In this research, we propose a simheuristic-based general variable neighbourhood search to solve the many-to-many hub location routing problem. This research contributes to the HLRP literature in multiple ways. There are not many publications on the hub location routing problem. The combination of hub locations, service areas and routing costs is relevant to be researched since these decisions are related to each other. By applying the HLRP to the network design of Company X, we expand the research on the HLRP. Also, apart from many authors in the HLRP field, our research was applied to a real-life (sized) case of Company X. The size of the case of Company X is large, presenting a challenge in the solving of the HLRP which is typically hard to solve because it is essentially a combination of two NP-Hard problems. Still, this research contributed to the formulation of models that are able to solve this problem for real-life instances. Additionally, to the best of our knowledge, the largest instance in the HLRP literature to date consisted of 2042 locations (Wasner & Zäpfel, 2004). The case in this research is considerably larger with a total of 4070 postal zones that are considered. On top of that, this research considers three logistical processes in contrast to the regular collection and delivery processes. Therefore, this research contributes to the literature by providing a method to solve HLRPs on a much larger instance. Lastly, this research is the first to propose a simheuristic approach for the HLRP. By doing so we created a model that is able to yield robust results given uncertain future scenarios and is also able to deal with HLRPs in a context of strong fluctuations in workload. Additionally, the number of authors that have included stochasticity in their HLRP formulations is to the best of our knowledge only one, and therefore this research is the second to include stochasticity in solving the HLRP.

Apart from the contributions to literature, we also provided a practical contribution to Company X. First of all, we determined an improved network with a cost reduction of 8.4%. Besides this, and maybe even of more value, the tool that we built in this research can be used by company X to achieve its ambition. The tool that was developed is a very detailed model, containing many elements related to the design of the logistical network. In the implementation, a lot of work is put into making sure the model is expandable and adaptable to future situations. Therefore, this tool enables Company X to do much more detailed network analyses in the future and will help them to achieve their goals of having a robust, efficient, and sustainable logistical network.

BIBLIOGRAPHY

Abbasi, M., Mokhtari, N., Shahvar, H., & Mahmoudi, A. (2019). Application of variable neighborhood search for solving large-scale many to many hub location routing problems. *Journal of Advances in Management Research*, *16*(5), 683–697. https://doi.org/10.1108/JAMR-11-2018-0107

Alumur, S. A., Campbell, J. F., Contreras, I., Kara, B. Y., Marianov, V., & O'Kelly, M. E. (2021). Perspectives on modeling hub location problems. *European Journal of Operational Research*, *291*(1), 1–17. https://doi.org/10.1016/j.ejor.2020.09.039

Alumur, S. A., Nickel, S., & Saldanha-da-Gama, F. (2012). Hub location under uncertainty. *Transportation Research Part B: Methodological*, *46*(4), 529–543. https://doi.org/10.1016/j.trb.2011.11.006

Basallo-Triana, M. J., Vidal-Holguín, C. J., & Bravo-Bastidas, J. J. (2021). Planning and design of intermodal hub networks: A literature review. *Computers & Operations Research*, *136*, 105469. https://doi.org/10.1016/j.cor.2021.105469

Basirati, M., Akbari Jokar, M. R., & Hassannayebi, E. (2020). Bi-objective optimization approaches to many-to-many hub location routing with distance balancing and hard time window. *Neural Computing and Applications*, *32*(17), 13267–13288. https://doi.org/10.1007/s00521-019-04666-z

Bostel, N., Dejax, P., & Mi Zhang. (2015). A model and a metaheuristic method for the Hub Location Routing Problem and application to postal services. *2015 International Conference on Industrial Engineering and Systems Management (IESM)*, 1383–1389. https://doi.org/10.1109/IESM.2015.7380332

Calvet, L., Wang, D., Juan, A., & Bové, L. (2019). Solving the multidepot vehicle routing problem with limited depot capacity and stochastic demands. *International Transactions in Operational Research*, *26*(2), 458–484. https://doi.org/10.1111/itor.12560

Campbell, J. F. (1994). Integer programming formulations of discrete hub location problems. *European Journal of Operational Research*, *72*(2), 387–405. https://doi.org/10.1016/0377-2217(94)90318-2

Catanzaro, D., Gourdin, E., Labbé, M., & Özsoy, F. A. (2011). A branch-and-cut algorithm for the partitioning-hub location-routing problem. *Computers & Operations Research*, *38*(2), 539–549. https://doi.org/10.1016/j.cor.2010.07.014

Çetiner, S., Sepil, C., & Süral, H. (2010). Hubbing and routing in postal delivery systems. *Annals of Operations Research*, *181*(1), 109–124. https://doi.org/10.1007/s10479-010-0705-2

Chica, M., A, J. A., Christopher, B., Oscar, C., & David, K. W. (2020). Why simheuristics? Benefits, limitations, and best practices when combining metaheuristics with simulation. *SORT-Statistics and Operations Research Transactions*, *44*(2). https://doi.org/10.2436/20.8080.02.104

Clarke, G., & Wright, J. W. (1964). Scheduling of Vehicles from a Central Depot to a Number of Delivery Points. *Operations Research*, *12*(4), 568–581. https://doi.org/10.1287/opre.12.4.568

Contreras, I. (2015). Hub Location Problems. In *Location Science* (pp. 311–344). Springer International Publishing. https://doi.org/10.1007/978-3-319-13111-5_12

Contreras, I. (2021). Hub Network Design. In *Network Design with Applications to Transportation and Logistics* (pp. 567–598). Springer International Publishing. https://doi.org/10.1007/978-3-030-64018-7_18

Crainic, T. G., Gendreau, M., & Gendron, B. (2021). *Network Design with Applications to Transportation and Logistics* (T. G. Crainic, M. Gendreau, & B. Gendron, Eds.). Springer International Publishing. https://doi.org/10.1007/978-3-030-64018-7

Crainic, T. G., & Hewitt, M. (2021). Service Network Design. In *Network Design with Applications to Transportation and Logistics* (pp. 347–382). Springer International Publishing. https://doi.org/10.1007/978-3-030-64018-7_12

Daganzo, C. F. (2005). *Logistics Systems Analysis*. Springer-Verlag. https://doi.org/10.1007/3-540-27516-9

Danach, K., Gelareh, S., & Neamatian Monemi, R. (2019). The capacitated single-allocation p-hub location routing problem: a Lagrangian relaxation and a hyper-heuristic approach. *EURO Journal on Transportation and Logistics*, *8*(5), 597–631. https://doi.org/10.1007/s13676-019-00141-w

de Armas, J., Juan, A. A., Marquès, J. M., & Pedroso, J. P. (2017). Solving the deterministic and stochastic uncapacitated facility location problem: from a heuristic to a simheuristic. *Journal of the Operational Research Society*, *68*(10), 1161–1176. https://doi.org/10.1057/s41274-016-0155-6 de Camargo, R. S., de Miranda, G., & Løkketangen, A. (2013). A new formulation and an exact approach for the many-to-many hub location-routing problem. *Applied Mathematical Modelling*, *37*(12–13), 7465–7480. https://doi.org/10.1016/j.apm.2013.02.035

Drexl, M., & Schneider, M. (2015). A survey of variants and extensions of the location-routing problem. *European Journal of Operational Research*, *241*(2), 283–308. https://doi.org/10.1016/J.EJOR.2014.08.030

Fontes, F. F. da C., & Goncalves, G. (2021). A variable neighbourhood decomposition search approach applied to a global liner shipping network using a hub-and-spoke with sub-hub structure. *International Journal of Production Research*, *59*(1), 30–46. https://doi.org/10.1080/00207543.2019.1693652

Gelareh, S., Monemic, R. N., & Semetd, F. (2015). *Capacitated Bounded Cardinality Hub Routing Problem: Model and Solution Algorithm*.

GeoDMS. (2019, January 29). *PC4 travel time matrix for Nederland van Boven*. https://www.geodms.nl/PC4_travel_time_matrix_for_Nederland_van_Boven

Ghaffarinasab, N., van Woensel, T., & Minner, S. (2018). A continuous approximation approach to the planar hub location-routing problem: Modeling and solution algorithms. *Computers & Operations Research*, *100*, 140–154. https://doi.org/10.1016/j.cor.2018.07.022

Gülpınar, N., Pachamanova, D., & Çanakoğlu, E. (2013). Robust strategies for facility location under uncertainty. *European Journal of Operational Research*, *225*(1), 21–35. https://doi.org/10.1016/j.ejor.2012.08.004

Habibzadeh Boukani, F., Farhang Moghaddam, B., & Pishvaee, M. S. (2016). Robust optimization approach to capacitated single and multiple allocation hub location problems. *Computational and Applied Mathematics*, *35*(1), 45–60. https://doi.org/10.1007/s40314-014-0179-y

Interprovinciaal Overleg. (2022, July 6). *IBIS Bedrijventerreinen*. https://data.overheid.nl/dataset/ibis-bedrijventerreinen

Juan, A. A., Faulin, J., Grasman, S. E., Rabe, M., & Figueira, G. (2015). A review of simheuristics: Extending metaheuristics to deal with stochastic combinatorial optimization problems. *Operations Research Perspectives*, *2*, 62–72. https://doi.org/10.1016/j.orp.2015.03.001

Karimi, H., & Setak, M. (2018). A bi-objective incomplete hub location-routing problem with flow shipment scheduling. *Applied Mathematical Modelling*, *57*, 406–431. https://doi.org/10.1016/J.APM.2018.01.012

Kartal, Z., Hasgul, S., & Ernst, A. T. (2017). Single allocation p-hub median location and routing problem with simultaneous pick-up and delivery. *Transportation Research Part E: Logistics and Transportation Review*, *108*, 141–159. https://doi.org/10.1016/j.tre.2017.10.004

Kartal, Z., Krishnamoorthy, M., & Ernst, A. T. (2019). Heuristic algorithms for the single allocation p-hub center problem with routing considerations. *OR Spectrum*, *41*(1), 99–145. https://doi.org/10.1007/s00291-018-0526-2

Khaleghi, A., & Eydi, A. (2022). Sustainable multi-period hub location under uncertainty. *Soft Computing*. https://doi.org/10.1007/s00500-022-07302-6

Law, A. M. (2014). *Simulation Modeling and Analysis* (Fifth). McGraw_Hill Education.

Lopes, M. C., de Andrade, C. E., de Queiroz, T. A., Resende, M. G. C., & Miyazawa, F. K. (2016). Heuristics for a hub location-routing problem. *Networks*, *68*(1), 54–90. https://doi.org/10.1002/net.21685

Magnanti, T. L., & Wong, R. T. (1984). Network Design and Transportation Planning: Models and Algorithms. *Transportation Science*, *18*(1), 1–55. https://doi.org/10.1287/trsc.18.1.1

Mara, S. T. W., Kuo, R. J., & Asih, A. M. S. (2021). Location-routing problem: a classification of recent research. *International Transactions in Operational Research*, *28*(6), 2941–2983. https://doi.org/10.1111/itor.12950

Mekamcha, K., Souier, M., Bessenouci, H. N., & Bennekrouf, M. (2021). Two metaheuristics approaches for solving the traveling salesman problem: an Algerian waste collection case. *Operational Research*, *21*(3), 1641–1661. https://doi.org/10.1007/s12351-019-00529-6

Min, H., Jayaraman, V., & Srivastava, R. (1998). Combined location-routing problems: A synthesis and future research directions. *European Journal of Operational Research*, *108*(1), 1–15. https://doi.org/10.1016/S0377-2217(97)00172-0

Mokhatari, N., & Abbasi, M. (2014). Applying VNPSO Algorithm to Solve the Many-to-Many Hub Location-Routing Problem in a Large scale. *European Online Journal of Natural and Social Sciences*, *3*(4), 647–656.

Nagy, G., & Salhi, S. (1998). The many-to-many location-routing problem. *Top*, *6*(2), 261–275. https://doi.org/10.1007/BF02564791

O'kelly, M. E. (1987). A quadratic integer program for the location of interacting hub facilities. *European Journal of Operational Research*, *32*(3), 393–404. https://doi.org/10.1016/S0377-2217(87)80007-3

Pandiri, V., & Singh, A. (2021). A simple hyper-heuristic approach for a variant of many-to-many hub location-routing problem. *Journal of Heuristics*, *27*(5), 791–868. https://doi.org/10.1007/s10732-021-09477-x

Postnord. (2021). *E-commerce in Europe 2021*. https://www.postnord.se/siteassets/pdf/rapporter/e-commerce-in-europe-2021.pdf

Ratli, M., Urošević, D., El Cadi, A. A., Brimberg, J., Mladenović, N., & Todosijević, R. (2022). An efficient heuristic for a hub location routing problem. *Optimization Letters*, *16*(1), 281–300. https://doi.org/10.1007/s11590-020-01675-z

Real, L. B., Contreras, I., Cordeau, J.-F., de Camargo, R. S., & de Miranda, G. (2021). Multimodal hub network design with flexible routes. *Transportation Research Part E: Logistics and Transportation Review*, *146*, 102188. https://doi.org/10.1016/j.tre.2020.102188

Rieck, J., Ehrenberg, C., & Zimmermann, J. (2014). Many-to-many location-routing with inter-hub transport and multi-commodity pickup-and-delivery. *European Journal of Operational Research*, *236*(3), 863–878. https://doi.org/10.1016/j.ejor.2013.12.021

Rodríguez-Martín, I., Salazar-González, J.-J., & Yaman, H. (2014). A branch-and-cut algorithm for the hub location and routing problem. *Computers & Operations Research*, *50*, 161–174. https://doi.org/10.1016/j.cor.2014.04.014

Ropke, S., & Pisinger, D. (2006). An Adaptive Large Neighborhood Search Heuristic for the Pickup and Delivery Problem with Time Windows. *Transportation Science*, *40*(4), 455–472. https://doi.org/10.1287/trsc.1050.0135

Rosenhead, J., Elton, M., & Gupta, S. K. (1972). Robustness and Optimality as Criteria for Strategic Decisions. *Operational Research Quarterly (1970-1977)*, *23*(4), 413. https://doi.org/10.2307/3007957

Seto, A., Uyama, K., Hosoda, J., & Miyashita, N. (2021). Hierarchical Clustering-Based Network Design Algorithm for Many-To-Many Hub Location Routing Problem. *2021 IEEE 17th International Conference on Automation Science and Engineering (CASE)*, 89–94. https://doi.org/10.1109/CASE49439.2021.9551550

Shahabi, M., & Unnikrishnan, A. (2014). Robust hub network design problem. *Transportation Research Part E: Logistics and Transportation Review*, *70*, 356–373. https://doi.org/10.1016/j.tre.2014.08.003

Shaw, P. (1998). *Using Constraint Programming and Local Search Methods to Solve Vehicle Routing Problems* (pp. 417–431). https://doi.org/10.1007/3-540-49481-2_30

Snyder, L. v. (2006). Facility location under uncertainty: a review. *IIE Transactions*, *38*(7), 547–564. https://doi.org/10.1080/07408170500216480

Statista. (2022a). *E-commerce in the Netherlands*. https://www-statista-com.ezproxy2.utwente.nl/study/38799/e-commerce-in-the-netherlands-statista-dossier/

Statista. (2022b). *The Netherlands' Economy & Society*. https://www-statista-com.ezproxy2.utwente.nl/study/48391/netherlands/

Sun, J. U. (2013). An Integrated Hub Location and Multi-Depot Vehicle Routing Problem. *Applied Mechanics and Materials*, *409–410*, 1188–1192. https://doi.org/10.4028/www.scientific.net/AMM.409-410.1188

Wasner, M., & Zäpfel, G. (2004). An integrated multi-depot hub-location vehicle routing model for network planning of parcel service. *International Journal of Production Economics*, *90*(3), 403–419. https://doi.org/10.1016/j.ijpe.2003.12.002

Watson-Gandy, C., & Dohrn, P. (1973). Depot location with van salesmen — A practical approach. *Omega*, *1*(3), 321–329. https://doi.org/10.1016/0305-0483(73)90108-4

Winkenbach, M., Kleindorfer, P. R., & Spinler, S. (2016). Enabling Urban Logistics Services at *La Poste* through Multi-Echelon Location-Routing. *Transportation Science*, *50*(2), 520–540. https://doi.org/10.1287/trsc.2015.0624

Wu, Y., Qureshi, A. G., & Yamada, T. (2022). Adaptive large neighborhood decomposition search algorithm for multi-allocation hub location routing problem. *European Journal of Operational Research*, *302*(3), 1113–1127. https://doi.org/10.1016/j.ejor.2022.02.002

Yahyaei, M., Bashiri, M., & Garmeyi, Y. (2014). Multicriteria Logistic Hub Location by Network Segmentation under Criteria Weights Uncertainty. *International Journal of Engineering*, *27*(8 (B)). https://doi.org/10.5829/idosi.ije.2014.27.08b.06

Yang, X., Bostel, N., & Dejax, P. (2019). A MILP model and memetic algorithm for the Hub Location and Routing problem with distinct collection and delivery tours. *Computers & Industrial Engineering*, *135*, 105–119. https://doi.org/10.1016/j.cie.2019.05.038

Zhou, L., Baldacci, R., Vigo, D., & Wang, X. (2018). A Multi-Depot Two-Echelon Vehicle Routing Problem with Delivery Options Arising in the Last Mile Distribution. *European Journal of Operational Research*, *265*(2), 765–778. https://doi.org/10.1016/j.ejor.2017.08.011

Zhou, L., Wang, X., Ni, L., & Lin, Y. (2016). Location-Routing Problem with Simultaneous Home Delivery and Customer's Pickup for City Distribution of Online Shopping Purchases. *Sustainability*, *8*(8), 828. https://doi.org/10.3390/su8080828

Appendix A. RESULTS OF EXPERIMENT 1

Table 13 Results of experiment 1

Run	$k_{\max}^{0}$	$k_{\max}^{1}$	Value	Run time (s)	Start hubs	Open hubs	Total route length	Total route duration	Total num. routes
1	3	3	7,688	86	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	2,170	114	15
2	3	3	6,341	47	[0, 2, 3, 7, 8, 9]	[0, 7, 8]	1,813	93	14
3	3	3	7,624	70	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	1,954	127	16
4	3	3	8,591	94	[0, 2, 3, 7, 8, 9]	[3, 4, 5, 7]	2,089	128	15
5	3	3	6,961	36	[0, 2, 3, 7, 8, 9]	[3, 7, 8, 9]	1,840	102	15
6	3	3	7,060	45	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	1,827	102	14
7	3	3	8,149	51	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	2,164	128	16
8	3	3	8,276	121	[0, 2, 3, 7, 8, 9]	[0, 6, 7]	2,250	125	17
9	3	3	7,144	45	[0, 2, 3, 7, 8, 9]	[2, 7, 8, 9]	1,618	92	13
10	3	3	6,952	138	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	1,929	117	15
11	3	5	7,632	131	[0, 2, 3, 7, 8, 9]	[0, 3, 8]	2,216	114	15
12	3	5	6,748	43	[0, 2, 3, 7, 8, 9]	[2, 3, 7, 8, 9]	1,668	90	14
13	3	5	7,600	197	[0, 2, 3, 7, 8, 9]	[0, 3, 8]	1,994	128	16
14	3	5	8,531	127	[0, 2, 3, 7, 8, 9]	[0, 2, 3, 6, 8]	1,927	125	16
15	3	5	7,933	62	[0, 2, 3, 7, 8, 9]	[0, 1, 3, 8, 9]	1,773	99	15
16	3	5	7,102	142	[0, 2, 3, 7, 8, 9]	[0, 3, 8]	1,939	106	14
17	3	5	8,378	135	[0, 2, 3, 7, 8, 9]	[0, 2, 3, 8]	2,080	127	16
18	3	5	8,036	252	[0, 2, 3, 7, 8, 9]	[0, 7, 8]	2,187	125	17
19	3	5	6,894	71	[0, 2, 3, 7, 8, 9]	[0, 3, 6]	1,743	97	13
20	3	5	6,953	86	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	1,929	118	15
21	3	8	7,686	330	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	2,180	114	15
22	3	8	6,946	131	[0, 2, 3, 7, 8, 9]	[0, 2, 3, 7, 8]	1,671	90	15
23	3	8	7,778	214	[0, 2, 3, 7, 8, 9]	[0, 7, 8]	2,056	129	16
24	3	8	7,911	184	[0, 2, 3, 7, 8, 9]	[0, 7, 8]	2,085	128	15
25	3	8	7,809	174	[0, 2, 3, 7, 8, 9]	[2, 3, 9]	2,069	106	15
26	3	8	7,014	290	[0, 2, 3, 7, 8, 9]	[0, 3, 6]	1,920	105	14
27	3	8	8,057	291	[0, 2, 3, 7, 8, 9]	[0, 1, 2, 3, 6]	1,851	118	16
28	3	8	8,314	102	[0, 2, 3, 7, 8, 9]	[0, 2, 7, 8]	2,077	122	17
29	3	8	6,942	162	[0, 2, 3, 7, 8, 9]	[0, 7, 8]	1,712	97	13
30	3	8	6,953	166	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	1,929	118	15
31	5	3	7,688	270	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	2,170	114	15
32	5	3	6,412	58	[0, 2, 3, 7, 8, 9]	[3, 7, 8, 9]	1,752	91	14
33	5	3	7,624	107	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	1,954	127	16
34	5	3	7,829	175	[0, 2, 3, 7, 8, 9]	[0, 3, 8]	2,073	128	15
35	5	3	6,898	122	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	1,925	104	15
36	5	3	7,172	57	[0, 2, 3, 7, 8, 9]	[0, 3, 8]	1,907	105	14
37	5	3	8,094	127	[0, 2, 3, 7, 8, 9]	[0, 2, 3, 6]	2,052	128	16
38	5	3	8,603	103	[0, 2, 3, 7, 8, 9]	[0, 1, 4, 7, 8]	1,950	121	17
39	5	3	6,894	99	[0, 2, 3, 7, 8, 9]	[0, 3, 8]	1,730	97	13
40	5	3	7,290	351	[0, 2, 3, 7, 8, 9]	[0, 3]	2,149	121	15
41	5	5	8,226	137	[0, 2, 3, 7, 8, 9]	[2, 3, 7, 8, 9]	1,985	110	15
42	5	5	6,624	158	[0, 2, 3, 7, 8, 9]	[0, 7, 8]	1,871	93	14
43	5	5	7,824	236	[0, 2, 3, 7, 8, 9]	[0, 7, 8]	2,051	129	16
44	5	5	7,881	247	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	2,042	127	15
45	5	5	6,897	157	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	1,925	104	15
46	5	5	7,060	468	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	1,827	102	14
47	5	5	7,930	315	[0, 2, 3, 7, 8, 9]	[2, 3, 6, 9]	2,010	126	16
48	5	5	8,509	186	[0, 2, 3, 7, 8, 9]	[1, 2, 4, 7]	2,141	123	17

Table 13 Results of experiment 1

Run	$k^0_{\max}$	$k^1_{\max}$	Value	Run time (s)	Start hubs	Open hubs	Total route length	Total route duration	Total num. routes
49	5	5	7,261	113	[0, 2, 3, 7, 8, 9]	[0, 7]	1,933	100	13
50	5	5	7,351	276	[0, 2, 3, 7, 8, 9]	[2, 3, 9]	1,994	118	15
51	5	8	7,634	346	[0, 2, 3, 7, 8, 9]	[0, 3, 8]	2,217	114	15
52	5	8	6,038	293	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	1,796	92	14
53	5	8	7,770	397	[0, 2, 3, 7, 8, 9]	[3, 4, 8]	2,011	127	16
54	5	8	7,948	440	[0, 2, 3, 7, 8, 9]	[1, 3, 4, 8]	1,864	123	15
55	5	8	6,961	137	[0, 2, 3, 7, 8, 9]	[3, 7, 8, 9]	1,840	102	15
56	5	8	7,060	235	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	1,827	102	14
57	5	8	8,134	243	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	2,130	127	16
58	5	8	7,785	351	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	2,125	124	17
59	5	8	7,144	67	[0, 2, 3, 7, 8, 9]	[2, 7, 8, 9]	1,618	92	13
60	5	8	7,394	217	[0, 2, 3, 7, 8, 9]	[2, 3, 9]	2,003	118	15
61	8	3	8,096	131	[0, 2, 3, 7, 8, 9]	[2, 6, 7, 9]	2,065	112	15
62	8	3	6,423	120	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	1,843	93	14
63	8	3	7,624	210	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	1,954	127	16
64	8	3	8,065	181	[0, 2, 3, 7, 8, 9]	[0, 6, 7, 8]	2,013	122	16
65	8	3	6,893	105	[0, 2, 3, 7, 8, 9]	[0, 3, 7, 8]	1,817	102	15
66	8	3	6,902	207	[0, 2, 3, 7, 8, 9]	[0, 6, 7]	1,985	107	15
67	8	3	8,031	190	[0, 2, 3, 7, 8, 9]	[2, 6, 7, 9]	1,997	127	16
68	8	3	7,736	254	[0, 2, 3, 7, 8, 9]	[0, 3, 8]	2,178	125	17
69	8	3	6,894	135	[0, 2, 3, 7, 8, 9]	[0, 3, 8]	1,730	97	13
70	8	3	7,089	202	[0, 2, 3, 7, 8, 9]	[0, 3, 8]	1,974	118	15
71	8	5	7,690	424	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	2,125	114	15
72	8	5	6,255	237	[0, 2, 3, 7, 8, 9]	[0, 7, 8]	1,815	90	14
73	8	5	7,845	218	[0, 2, 3, 7, 8, 9]	[3, 7, 8, 9]	1,891	126	16
74	8	5	8,337	166	[0, 2, 3, 7, 8, 9]	[2, 3, 7, 8, 9]	1,893	124	16
75	8	5	6,897	288	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	1,925	104	15
76	8	5	7,128	319	[0, 2, 3, 7, 8, 9]	[0, 2, 7]	2,041	107	15
77	8	5	8,157	259	[0, 2, 3, 7, 8, 9]	[2, 3, 6, 7, 9]	1,932	124	16
78	8	5	8,035	461	[0, 2, 3, 7, 8, 9]	[0, 7, 8]	2,184	125	17
79	8	5	7,296	146	[0, 2, 3, 7, 8, 9]	[7, 8, 9]	1,747	97	13
80	8	5	6,952	245	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	1,929	117	15
81	8	8	7,652	483	[0, 2, 3, 7, 8, 9]	[0, 3, 8]	2,164	113	15
82	8	8	6,387	627	[0, 2, 3, 7, 8, 9]	[0, 1, 3, 8]	1,674	90	14
83	8	8	7,822	469	[0, 2, 3, 7, 8, 9]	[0, 7, 8]	2,049	129	16
84	8	8	7,910	360	[0, 2, 3, 7, 8, 9]	[0, 7, 8]	2,084	128	15
85	8	8	6,898	228	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	1,925	104	15
86	8	8	7,360	424	[0, 2, 3, 7, 8, 9]	[0, 3]	2,205	110	15
87	8	8	8,366	206	[0, 2, 3, 7, 8, 9]	[2, 7, 8, 9]	2,057	128	16
88	8	8	7,977	606	[0, 2, 3, 7, 8, 9]	[0, 7, 8]	2,185	125	17
89	8	8	7,466	212	[0, 2, 3, 7, 8, 9]	[0, 7, 8, 9]	1,634	95	13
90	8	8	7,708	316	[0, 2, 3, 7, 8, 9]	[2, 7, 9]	2,036	119	15

Appendix B. RESULTS OF EXPERIMENT 2

Table 14 Results of experiment 2: Operator performance

Run	Operator	Value	Run time (s)	Run	Operator	Value	Run time (s)
1	None	8,096	131	46	move_zone_PC2	7,172	86
2	None	6,423	120	47	move_zone_PC2	8,053	118
3	None	7,624	210	48	move_zone_PC2	7,733	183
4	None	8,065	181	49	move_zone_PC2	6,894	67
5	None	6,893	105	50	move_zone_PC2	7,089	235
6	None	6,902	207	51	move_zone_PC3	7,956	127
7	None	8,031	190	52	move_zone_PC3	6,388	54
8	None	7,736	254	53	move_zone_PC3	7,850	81
9	None	6,894	135	54	move_zone_PC3	7,880	154
10	None	7,089	202	55	move_zone_PC3	7,237	65
11	add_hub	8,607	171	56	move_zone_PC3	7,060	80
12	add_hub	6,830	67	57	move_zone_PC3	8,049	214
13	add_hub	8,224	149	58	move_zone_PC3	7,715	204
14	add_hub	7,884	146	59	move_zone_PC3	7,374	56
15	add_hub	6,893	77	60	move_zone_PC3	7,212	121
16	add_hub	7,337	115	61	swap_zones	8,311	311
17	add_hub	8,091	96	62	swap_zones	6,342	83
18	add_hub	8,043	105	63	swap_zones	8,147	90
19	add_hub	6,942	81	64	swap_zones	7,828	107
20	add_hub	6,952	137	65	swap_zones	7,298	53
21	remove_hub	8,887	18	66	swap_zones	7,604	68
22	remove_hub	7,062	16	67	swap_zones	8,784	147
23	remove_hub	8,603	17	68	swap_zones	8,037	198
24	remove_hub	8,899	14	69	swap_zones	6,838	96
25	remove_hub	7,952	13	70	swap_zones	7,712	114
26	remove_hub	8,240	21	71	move_zone_to_hub_Allp	7,903	210
27	remove_hub	8,865	23	72	move_zone_to_hub_Allp	6,231	92
28	remove_hub	8,525	30	73	move_zone_to_hub_Allp	8,272	111
29	remove_hub	7,866	11	74	move_zone_to_hub_Allp	8,468	82
30	remove_hub	8,505	17	75	move_zone_to_hub_Allp	6,915	198
31	swap_hubs	8,134	77	76	move_zone_to_hub_Allp	7,298	71
32	swap_hubs	6,398	37	77	move_zone_to_hub_Allp	7,985	151
33	swap_hubs	7,624	152	78	move_zone_to_hub_Allp	8,037	147
34	swap_hubs	8,254	57	79	move_zone_to_hub_Allp	6,963	60
35	swap_hubs	7,616	77	80	move_zone_to_hub_Allp	7,248	169
36	swap_hubs	7,258	73	81	move_zone_to_hub_1p	7,879	151
37	swap_hubs	8,537	75	82	move_zone_to_hub_1p	6,748	47
38	swap_hubs	7,918	210	83	move_zone_to_hub_1p	7,598	130
39	swap_hubs	6,974	73	84	move_zone_to_hub_1p	7,829	105
40	swap_hubs	7,629	124	85	move_zone_to_hub_1p	7,259	82
41	move_zone_PC2	7,653	139	86	move_zone_to_hub_1p	7,226	100
42	move_zone_PC2	6,694	91	87	move_zone_to_hub_1p	8,186	172
43	move_zone_PC2	7,724	116	88	move_zone_to_hub_1p	8,081	78
44	move_zone_PC2	7,836	68	89	move_zone_to_hub_1p	6,975	76
45	move_zone_PC2	7,473	75	90	move_zone_to_hub_1p	7,090	98

Table 15 Results of experiment 2: Operator order

Run	Operator Order	Value	Run time (s)	start hubs	open hubs	Total route length	Total route duration	Total num. routes
1	AH-RH-SH	7,652	277	[0, 2, 3, 7, 8, 9]	[0, 3, 8]	2,164	113	15
2	AH-RH-SH	6,051	229	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	1,803	92	14
3	AH-RH-SH	7,623	121	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	1,954	127	16
4	AH-RH-SH	7,881	337	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	2,042	127	15
5	AH-RH-SH	6,893	88	[0, 2, 3, 7, 8, 9]	[0, 3, 7, 8]	1,817	102	15
6	AH-RH-SH	7,015	165	[0, 2, 3, 7, 8, 9]	[0, 3, 6]	1,920	106	14
7	AH-RH-SH	8,091	135	[0, 2, 3, 7, 8, 9]	[0, 3, 8]	2,147	128	16
8	AH-RH-SH	7,733	242	[0, 2, 3, 7, 8, 9]	[0, 3, 8]	2,173	125	17
9	AH-RH-SH	6,894	166	[0, 2, 3, 7, 8, 9]	[0, 3, 8]	1,730	97	13
10	AH-RH-SH	7,089	121	[0, 2, 3, 7, 8, 9]	[0, 3, 8]	1,975	118	15
11	AH-SH-RH	7,990	248	[0, 2, 3, 7, 8, 9]	[0, 3, 6]	2,309	117	15
12	AH-SH-RH	6,749	114	[0, 2, 3, 7, 8, 9]	[1, 3, 8]	1,865	94	14
13	AH-SH-RH	8,007	152	[0, 2, 3, 7, 8, 9]	[3, 6, 8, 9]	1,897	126	16
14	AH-SH-RH	8,022	145	[0, 2, 3, 7, 8, 9]	[3, 4, 5, 8]	1,930	124	15
15	AH-SH-RH	7,238	107	[0, 2, 3, 7, 8, 9]	[0, 7, 8]	1,959	104	15
16	AH-SH-RH	7,172	137	[0, 2, 3, 7, 8, 9]	[0, 3, 8]	1,907	105	14
17	AH-SH-RH	8,803	288	[0, 2, 3, 7, 8, 9]	[0, 1, 3]	2,262	125	16
18	AH-SH-RH	8,037	198	[0, 2, 3, 7, 8, 9]	[0, 7, 8]	2,187	125	17
19	AH-SH-RH	6,871	128	[0, 2, 3, 7, 8, 9]	[0, 3, 8]	1,743	95	13
20	AH-SH-RH	6,947	77	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	1,927	117	15
21	RH-AH-SH	7,709	101	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	2,170	114	15
22	RH-AH-SH	6,398	53	[0, 2, 3, 7, 8, 9]	[3, 7, 8, 9]	1,745	91	14
23	RH-AH-SH	7,624	86	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	1,954	127	16
24	RH-AH-SH	8,334	71	[0, 2, 3, 7, 8, 9]	[2, 7, 8, 9]	2,028	127	15
25	RH-AH-SH	6,893	50	[0, 2, 3, 7, 8, 9]	[0, 3, 7, 8]	1,817	102	15
26	RH-AH-SH	6,912	162	[0, 2, 3, 7, 8, 9]	[0, 6, 7]	1,939	106	14
27	RH-AH-SH	8,087	181	[0, 2, 3, 7, 8, 9]	[0, 3, 8]	2,148	128	16
28	RH-AH-SH	8,023	104	[0, 2, 3, 7, 8, 9]	[2, 3, 8, 9]	2,005	121	17
29	RH-AH-SH	7,261	79	[0, 2, 3, 7, 8, 9]	[0, 7]	1,933	101	13
30	RH-AH-SH	7,284	217	[0, 2, 3, 7, 8, 9]	[3, 6, 8, 9]	1,877	113	15
31	RH-SH-AH	7,633	100	[0, 2, 3, 7, 8, 9]	[0, 3, 8]	2,217	114	15
32	RH-SH-AH	6,052	78	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	1,803	92	14
33	RH-SH-AH	8,148	85	[0, 2, 3, 7, 8, 9]	[0, 2, 7, 8]	1,977	127	16
34	RH-SH-AH	8,012	100	[0, 2, 3, 7, 8, 9]	[0, 3, 6, 8]	1,995	121	16
35	RH-SH-AH	7,328	94	[0, 2, 3, 7, 8, 9]	[2, 3, 8, 9]	1,862	103	15
36	RH-SH-AH	7,219	130	[0, 2, 3, 7, 8, 9]	[2, 7, 9]	1,951	106	15
37	RH-SH-AH	8,294	81	[0, 2, 3, 7, 8, 9]	[0, 2, 7, 8]	2,052	127	16
38	RH-SH-AH	8,035	118	[0, 2, 3, 7, 8, 9]	[0, 7, 8]	2,184	125	17
39	RH-SH-AH	7,027	95	[0, 2, 3, 7, 8, 9]	[3, 7, 8, 9]	1,606	95	13
40	RH-SH-AH	7,338	199	[0, 2, 3, 7, 8, 9]	[7, 8, 9]	1,946	117	15
41	SH-AH-RH	7,634	132	[0, 2, 3, 7, 8, 9]	[0, 3, 8]	2,217	114	15
42	SH-AH-RH	6,556	59	[0, 2, 3, 7, 8, 9]	[0, 2, 3, 8]	1,721	91	14
43	SH-AH-RH	8,101	118	[0, 2, 3, 7, 8, 9]	[1, 2, 3, 4]	2,030	128	16
44	SH-AH-RH	7,828	89	[0, 2, 3, 7, 8, 9]	[0, 3, 8]	2,073	128	15
45	SH-AH-RH	7,328	96	[0, 2, 3, 7, 8, 9]	[2, 3, 8, 9]	1,862	103	15
46	SH-AH-RH	7,243	106	[0, 2, 3, 7, 8, 9]	[0, 3, 8]	1,929	106	15
47	SH-AH-RH	8,091	179	[0, 2, 3, 7, 8, 9]	[0, 3, 8]	2,147	128	16
48	SH-AH-RH	8,082	152	[0, 2, 3, 7, 8, 9]	[3, 7, 8, 9]	2,047	123	17
49	SH-AH-RH	6,934	109	[0, 2, 3, 7, 8, 9]	[0, 7, 8]	1,719	96	13
50	SH-AH-RH	7,077	251	[0, 2, 3, 7, 8, 9]	[0, 3, 8]	1,972	118	15

Table 15 Results of experiment 2: Operator order

Run	Operator Order	Value	Run time (s)	start hubs	open hubs	Total route length	Total route duration	Total num. routes
51	SH-RH-AH	7,936	68	[0, 2, 3, 7, 8, 9]	[3, 7, 8, 9]	2,044	112	15
52	SH-RH-AH	6,231	84	[0, 2, 3, 7, 8, 9]	[0, 3, 8]	1,803	93	14
53	SH-RH-AH	7,599	106	[0, 2, 3, 7, 8, 9]	[0, 3, 8]	1,994	128	16
54	SH-RH-AH	7,829	77	[0, 2, 3, 7, 8, 9]	[0, 3, 8]	2,073	128	15
55	SH-RH-AH	7,410	57	[0, 2, 3, 7, 8, 9]	[0, 3, 8, 9]	1,833	101	15
56	SH-RH-AH	7,085	121	[0, 2, 3, 7, 8, 9]	[2, 6, 7, 9]	1,749	102	15
57	SH-RH-AH	8,287	133	[0, 2, 3, 7, 8, 9]	[2, 3, 8, 9]	2,041	128	16
58	SH-RH-AH	8,036	148	[0, 2, 3, 7, 8, 9]	[0, 7, 8]	2,187	125	17
59	SH-RH-AH	7,301	123	[0, 2, 3, 7, 8, 9]	[2, 3, 6, 9]	1,676	96	13
60	SH-RH-AH	6,952	154	[0, 2, 3, 7, 8, 9]	[3, 8, 9]	1,929	117	15

Appendix C. RESULTS OF EXPERIMENT 3

Table 16 Results of experiment 3

Run	Procedure	Value	Run time (s)	Start hubs	Open hubs	Total route length	Total route duration	Total num. routes
uu1	Random	8,095	195	[5, 7, 1, 3, 4]	[0, 1, 3, 8]	2,079	112	15
2	Random	6,535	177	[9, 0, 4]	[0, 1, 7, 8]	1,688	91	14
3	Random	7,599	118	[5, 7, 9, 3, 6, 2, 8, 4, 0]	[0, 3, 8]	1,994	128	16
4	Random	8,704	68	[9, 2, 5, 1, 4, 0, 8, 3, 7]	[1, 2, 3, 6, 9]	1,895	125	16
5	Random	6,961	89	[9, 2, 4, 7, 8, 6]	[3, 7, 8, 9]	1,840	102	15
6	Random	7,163	106	[3, 1, 6, 0, 4, 2, 8, 7]	[3, 6, 8, 9]	1,738	101	14
7	Random	8,363	218	[1, 9]	[2, 7, 8, 9]	2,059	128	16
8	Random	7,723	322	[9, 8, 5, 4, 6, 1]	[0, 3, 8]	2,173	125	17
9	Random	6,894	115	[0, 9, 8, 4, 3]	[0, 3, 6]	1,743	97	13
10	Random	7,284	121	[3, 1, 6, 0, 8, 7, 5]	[3, 6, 8, 9]	1,875	113	15
11	GrdCls_k4	7,634	142	[0, 3, 7, 8]	[0, 3, 8]	2,217	114	15
12	GrdCls_k4	6,513	138	[0, 3, 7, 8]	[1, 3, 8, 9]	1,720	91	14
13	GrdCls_k4	7,624	251	[0, 3, 7, 8]	[3, 8, 9]	1,954	127	16
14	GrdCls_k4	7,829	172	[0, 3, 7, 8]	[0, 3, 8]	2,073	128	15
15	GrdCls_k4	6,893	85	[0, 3, 7, 8]	[0, 3, 7, 8]	1,817	102	15
16	GrdCls_k4	6,983	74	[0, 2, 6, 7]	[0, 2, 6, 7]	1,819	104	15
17	GrdCls_k4	7,945	231	[0, 2, 6, 7]	[2, 3, 6, 9]	1,973	125	16
18	GrdCls_k4	8,029	110	[0, 3, 7, 8]	[0, 3, 7, 8]	2,095	123	17
19	GrdCls_k4	6,878	107	[0, 3, 6, 7]	[0, 3, 8]	1,746	95	13
20	GrdCls_k4	7,338	78	[0, 3, 7, 8]	[7, 8, 9]	1,946	117	15
21	GrdCls_k5	7,835	122	[0, 3, 4, 7, 8]	[0, 3, 4, 8]	2,110	108	15
22	GrdCls_k5	6,586	101	[0, 1, 3, 7, 8]	[0, 3]	2,024	95	14
23	GrdCls_k5	7,822	179	[0, 3, 6, 7, 8]	[0, 7, 8]	2,049	129	16
24	GrdCls_k5	8,095	134	[0, 3, 4, 7, 8]	[0, 6, 7, 8]	2,006	121	16
25	GrdCls_k5	7,371	86	[0, 3, 6, 7, 8]	[0, 3, 5, 8]	1,877	103	15
26	GrdCls_k5	7,014	159	[0, 2, 3, 6, 7]	[0, 3, 6]	1,920	106	14
27	GrdCls_k5	8,133	242	[0, 2, 3, 6, 7]	[3, 8, 9]	2,130	127	16
28	GrdCls_k5	7,945	201	[0, 3, 6, 7, 8]	[0, 3, 6, 8]	2,100	123	17
29	GrdCls_k5	7,180	110	[0, 1, 3, 6, 7]	[0, 3]	1,939	100	13
30	GrdCls_k5	7,090	135	[0, 3, 4, 7, 8]	[0, 3, 8]	1,974	118	15
31	GrdCls_k6	7,861	133	[0, 2, 3, 4, 7, 8]	[3, 4, 5, 8]	2,112	107	15
32	GrdCls_k6	6,760	172	[0, 1, 2, 3, 7, 8]	[2, 3, 6, 9]	1,796	93	14
33	GrdCls_k6	7,599	242	[0, 2, 3, 6, 7, 8]	[0, 3, 8]	1,994	128	16
34	GrdCls_k6	8,271	87	[0, 3, 4, 6, 7, 8]	[3, 4, 7, 8]	1,990	126	15
35	GrdCls_k6	7,193	164	[0, 3, 4, 6, 7, 8]	[3, 4, 8]	2,025	106	15
36	GrdCls_k6	7,172	112	[0, 1, 2, 3, 6, 7]	[0, 3, 8]	1,907	105	14
37	GrdCls_k6	8,197	113	[0, 1, 2, 3, 6, 7]	[1, 2, 4, 6, 7]	1,911	126	16
38	GrdCls_k6	7,768	177	[0, 1, 3, 6, 7, 8]	[3, 8, 9]	2,113	123	17
39	GrdCls_k6	6,893	82	[0, 1, 2, 3, 6, 7]	[0, 3, 6]	1,743	97	13
40	GrdCls_k6	6,952	184	[0, 3, 4, 6, 7, 8]	[3, 8, 9]	1,929	117	15
41	GrdCls_k7	7,926	200	[0, 2, 3, 4, 5, 7, 8]	[3, 8, 9]	2,249	115	15
42	GrdCls_k7	6,231	104	[0, 1, 2, 3, 4, 7, 8]	[0, 3, 8]	1,803	93	14
43	GrdCls_k7	7,822	224	[0, 2, 3, 4, 6, 7, 8]	[0, 7, 8]	2,049	129	16
44	GrdCls_k7	8,401	62	[0, 1, 3, 4, 6, 7, 8]	[1, 3, 4, 6, 7, 8]	1,717	121	16
45	GrdCls_k7	7,642	66	[0, 2, 3, 4, 6, 7, 8]	[1, 2, 3, 4, 7, 8]	1,693	99	15
46	GrdCls_k7	7,234	106	[0, 1, 2, 3, 4, 6, 7]	[0, 1, 6, 7]	1,843	104	14
47	GrdCls_k7	8,133	240	[0, 1, 2, 3, 4, 6, 7]	[0, 3, 6]	2,182	130	16
48	GrdCls_k7	8,141	151	[0, 1, 2, 3, 6, 7, 8]	[0, 1, 3, 6, 7]	1,948	120	17

Table 16 Results of experiment 3

Run	Procedure	Value	Run time (s)	Start hubs	Open hubs	Total route length	Total route duration	Total num. routes
49	GrdCls_k7	7,251	95	[0, 1, 2, 3, 4, 6, 7]	[0, 1, 3, 8]	1,629	95	13
50	GrdCls_k7	7,151	132	[0, 1, 3, 4, 6, 7, 8]	[0, 3, 8]	1,991	118	15
51	GrdCls_k8	7,690	146	[0, 2, 3, 4, 5, 6, 7, 8]	[3, 8, 9]	2,125	114	15
52	GrdCls_k8	6,608	389	[0, 1, 2, 3, 4, 6, 7, 8]	[0, 3]	2,009	94	14
53	GrdCls_k8	7,863	200	[0, 1, 2, 3, 4, 6, 7, 8]	[1, 3, 4, 8]	1,878	126	16
54	GrdCls_k8	7,881	200	[0, 1, 2, 3, 4, 6, 7, 8]	[3, 8, 9]	2,042	127	15
55	GrdCls_k8	7,240	98	[0, 2, 3, 4, 5, 6, 7, 8]	[0, 3, 4, 8]	1,866	103	15
56	GrdCls_k8	7,172	206	[0, 1, 2, 3, 4, 6, 7, 9]	[0, 3, 8]	1,907	105	14
57	GrdCls_k8	8,205	158	[0, 1, 2, 3, 4, 5, 6, 7]	[3, 6, 8, 9]	2,042	126	16
58	GrdCls_k8	8,274	95	[0, 1, 2, 3, 4, 6, 7, 8]	[0, 5, 6, 7]	2,046	121	17
59	GrdCls_k8	7,296	64	[0, 1, 2, 3, 4, 6, 7, 9]	[7, 8, 9]	1,747	97	13
60	GrdCls_k8	7,307	193	[0, 1, 2, 3, 4, 6, 7, 8]	[0, 7, 8]	1,997	119	15
61	ImprCls_k4	7,890	141	[0, 3, 7, 8]	[0, 3, 7, 8]	2,103	112	15
62	ImprCls_k4	6,342	84	[0, 3, 7, 8]	[0, 7, 8]	1,813	93	14
63	ImprCls_k4	7,624	130	[0, 3, 7, 8]	[3, 8, 9]	1,954	127	16
64	ImprCls_k4	7,909	203	[0, 3, 7, 8]	[0, 7, 8]	2,081	128	15
65	ImprCls_k4	6,830	104	[0, 3, 7, 8]	[0, 3, 8]	1,902	103	15
66	ImprCls_k4	6,983	123	[0, 2, 6, 7]	[0, 2, 6, 7]	1,819	104	15
67	ImprCls_k4	8,065	90	[0, 2, 6, 7]	[0, 2, 6, 7]	2,029	127	16
68	ImprCls_k4	7,877	179	[0, 3, 7, 8]	[0, 3, 6]	2,228	124	17
69	ImprCls_k4	6,894	91	[0, 3, 6, 7]	[0, 3, 6]	1,743	97	13
70	ImprCls_k4	7,336	117	[0, 3, 7, 8]	[0, 3, 4, 8]	1,888	117	16
71	ImprCls_k5	8,060	268	[0, 3, 4, 7, 8]	[0, 3, 6]	2,263	116	16
72	ImprCls_k5	6,342	105	[0, 1, 3, 7, 8]	[0, 7, 8]	1,813	93	14
73	ImprCls_k5	7,924	147	[0, 3, 6, 7, 8]	[0, 1, 3, 8]	1,891	126	16
74	ImprCls_k5	8,291	199	[0, 3, 4, 7, 8]	[0, 3, 6]	2,093	128	15
75	ImprCls_k5	7,375	62	[0, 3, 6, 7, 8]	[0, 3, 6, 7, 8]	1,780	101	15
76	ImprCls_k5	6,908	174	[0, 2, 3, 6, 7]	[0, 6, 7]	1,998	107	15
77	ImprCls_k5	8,031	137	[0, 2, 3, 6, 7]	[2, 6, 7, 9]	1,997	127	16
78	ImprCls_k5	7,701	362	[0, 3, 6, 7, 8]	[0, 3, 8]	2,186	125	17
79	ImprCls_k5	7,256	44	[0, 1, 3, 6, 7]	[0, 1, 6, 7]	1,640	95	13
80	ImprCls_k5	7,429	181	[0, 3, 4, 7, 8]	[0, 6, 7]	2,030	119	15
81	ImprCls_k6	7,686	218	[0, 2, 3, 4, 7, 8]	[3, 8, 9]	2,180	114	15
82	ImprCls_k6	6,231	160	[0, 1, 2, 3, 7, 8]	[0, 3, 8]	1,803	93	14
83	ImprCls_k6	7,599	133	[0, 2, 3, 6, 7, 8]	[0, 3, 8]	1,993	128	16
84	ImprCls_k6	7,828	119	[0, 3, 4, 6, 7, 8]	[0, 3, 8]	2,073	128	15
85	ImprCls_k6	6,830	117	[0, 3, 4, 6, 7, 8]	[0, 3, 8]	1,902	103	15
86	ImprCls_k6	6,913	193	[0, 1, 2, 3, 6, 7]	[0, 6, 7]	1,939	106	14
87	ImprCls_k6	8,268	49	[0, 1, 2, 3, 6, 7]	[0, 1, 2, 3, 6, 7]	1,809	118	16
88	ImprCls_k6	7,993	313	[0, 1, 3, 6, 7, 8]	[0, 3, 5, 8]	2,028	122	18
89	ImprCls_k6	7,180	99	[0, 1, 2, 3, 6, 7]	[0, 3]	1,939	100	13
90	ImprCls_k6	7,290	109	[0, 3, 4, 6, 7, 8]	[3, 4, 5, 8]	1,935	117	16
91	ImprCls_k7	8,159	186	[0, 2, 3, 4, 5, 7, 8]	[0, 3, 8, 9]	2,062	112	15
92	ImprCls_k7	6,301	178	[0, 1, 2, 3, 4, 7, 8]	[0, 3, 6]	1,895	94	14
93	ImprCls_k7	7,585	234	[0, 2, 3, 4, 6, 7, 8]	[0, 3, 8]	1,995	127	16
94	ImprCls_k7	7,948	300	[0, 1, 3, 4, 6, 7, 8]	[1, 3, 4, 8]	1,864	123	15
95	ImprCls_k7	6,898	130	[0, 2, 3, 4, 6, 7, 8]	[3, 8, 9]	1,925	104	15
96	ImprCls_k7	6,990	182	[0, 1, 2, 3, 4, 6, 7]	[0, 6, 7]	2,007	107	15
97	ImprCls_k7	8,039	115	[0, 1, 2, 3, 4, 6, 7]	[0, 1, 2, 3, 6]	1,892	119	16
98	ImprCls_k7	8,023	179	[0, 1, 2, 3, 6, 7, 8]	[2, 3, 8, 9]	2,005	121	17
99	ImprCls_k7	6,916	232	[0, 1, 2, 3, 4, 6, 7]	[0, 6, 7]	1,729	97	13

Table 16 Results of experiment 3

Run	Procedure	Value	Run time (s)	Start hubs	Open hubs	Total route length	Total route duration	Total num. routes
100	ImprCls_k7	7,543	219	[0, 1, 3, 4, 6, 7, 8]	[1, 2, 4, 7]	1,880	116	16
101	ImprCls_k8	7,922	233	[0, 2, 3, 4, 5, 6, 7, 8]	[3, 4, 8]	2,250	111	15
102	ImprCls_k8	6,342	155	[0, 1, 2, 3, 4, 6, 7, 8]	[0, 7, 8]	1,813	93	14
103	ImprCls_k8	7,823	174	[0, 1, 2, 3, 4, 6, 7, 8]	[0, 7, 8]	2,051	129	16
104	ImprCls_k8	7,881	125	[0, 1, 2, 3, 4, 6, 7, 8]	[3, 8, 9]	2,042	127	15
105	ImprCls_k8	7,193	180	[0, 2, 3, 4, 5, 6, 7, 8]	[3, 4, 8]	2,025	106	15
106	ImprCls_k8	7,441	203	[0, 1, 2, 3, 4, 6, 7, 9]	[0, 8]	2,285	112	15
107	ImprCls_k8	8,545	178	[0, 1, 2, 3, 4, 5, 6, 7]	[0, 2, 5, 7]	2,149	129	16
108	ImprCls_k8	8,133	109	[0, 1, 2, 3, 4, 6, 7, 8]	[0, 1, 6, 7]	2,016	120	17
109	ImprCls_k8	7,292	174	[0, 1, 2, 3, 4, 6, 7, 9]	[3, 4, 5, 8]	1,709	97	13
110	ImprCls_k8	6,947	170	[0, 1, 2, 3, 4, 6, 7, 8]	[3, 8, 9]	1,927	117	15

Appendix D. RESULTS OF EXPERIMENTS 6, 7, AND 8

Table 17 Results of experiments 6, 7, 8

	Solution		
	Simheuristic GVNS Experiment 6	**Simple method Experiment 7**	**Initial Plan Experiment 8**
Hub costs	6,061	5,455	6,061
Linehaul costs	5,553	5,550	5,420
Routing costs	12,614	12,656	12,582
DO-point cost	727	727	727
Penalty cost	334	351	2,471
Total	25,271	24,739	27,261
Num. hubs	2	1	2
Num. routes	83	83	84
Total route km	8,448	8,521	8,354
Total route duration	320	322	319
Hub costs	6,061	5,455	6,061

Appendix E. RESULTS OF EXPERIMENT 8

The contents of this appendix have been shielded because of their relevance to the company.

Shielded for the public version

Solution 1

Average	€ 27,261.18
Standard Deviation	€ 5,606.83

Solution 2

Average	€ 27,263.98
Standard Deviation	€ 5,609.16

Shielded for the public version

Solution 3

Average	€ 27,277.03
Standard Deviation	€ 5,611.75

Solution 4

Average	€ 27,296.48
Standard Deviation	€ 5,624.62

Shielded for the public version

Solution 5

Average	€ 27,296.48
Standard Deviation	€ 5,624.63

Solution 6

Average	€ 27,296.67
Standard Deviation	€ 5,625.05

Appendix F. RESULTS OF EXPERIMENT 9

Table 18 Results of experiment 9: Sensitivity analysis - objective

		Objective				
Exp	**Input parameter**	**50%**	**75%**	**100%**	**125%**	**150%**
9.1	Hub opening cost	6,692	7,255	7,861	8,444	8,941
9.2	PPC duration	8,139	8,060	7,956	7,513	8,248
9.3	HD duration	8,424	7,854	8,189	7,579	7,643
9.4	Fixed vehicle cost	7,344	8,129	7,956	7,786	7,937
9.5	Hourly vehicle cost	6,458	7,158	7,689	8,283	9,040
9.6	Cost per km	7,678	7,796	7,825	7,866	7,691
9.7	Vehicle capacity	7,876	8,228	7,634	7,894	7,989
9.8	Vehicle driving range	8,132	8,188	7,642	7,560	7,711
9.9	Linehaul cost factor	7,685	7,688	7,956	7,602	7,964

Table 19 Results of experiment 9: Sensitivity analysis – opened hubs[8]

		Number of opened hubs				
Exp	**Input parameter**	**50%**	**75%**	**100%**	**125%**	**150%**
9.1	Hub opening cost	1.70	1.13	1.13	0.85	0.85
9.2	PPC duration	1.13	1.13	1.13	0.85	1.42
9.3	HD duration	0.85	0.85	1.13	0.85	1.13
9.4	Fixed vehicle cost	0.85	0.85	1.13	0.85	0.85
9.5	Hourly vehicle cost	1.13	1.13	0.85	0.85	0.85
9.6	Cost per km	1.13	0.85	1.13	1.13	0.85
9.7	Vehicle capacity	0.85	1.13	0.85	1.13	0.85
9.8	Vehicle driving range	0.85	1.13	0.85	0.85	1.13
9.9	Linehaul cost factor	0.85	0.85	1.13	0.85	1.13

Table 20 Results of experiment 9: Sensitivity analysis –Hub opening costs

		Hub opening costs				
Exp	**Input parameter**	**50%**	**75%**	**100%**	**125%**	**150%**
9.1	Hub opening cost	1,818	1,818	2,424	2,273	2,727
9.2	PPC duration	2,424	2,424	2,424	1,818	3,030
9.3	HD duration	1,818	1,818	2,424	1,818	2,424
9.4	Fixed vehicle cost	1,818	1,818	2,424	1,818	1,818
9.5	Hourly vehicle cost	2,424	2,424	1,818	1,818	1,818
9.6	Cost per km	2,424	1,818	2,424	2,424	1,818
9.7	Vehicle capacity	1,818	2,424	1,818	2,424	1,818
9.8	Vehicle driving range	1,818	2,424	1,818	1,818	2,424
9.9	Linehaul cost factor	1,818	1,818	2,424	1,818	2,424

[8] *Relative to the average*

Table 21 Results of experiment 9: Sensitivity analysis – Linehaul costs

		Linehaul costs				
Exp	**Input parameter**	**50%**	**75%**	**100%**	**125%**	**150%**
9.1	Hub opening cost	40	30	30	38	28
9.2	PPC duration	38	37	39	37	31
9.3	HD duration	27	37	38	28	37
9.4	Fixed vehicle cost	24	31	39	24	24
9.5	Hourly vehicle cost	30	38	37	37	37
9.6	Cost per km	29	27	38	26	24
9.7	Vehicle capacity	28	38	24	37	38
9.8	Vehicle driving range	37	29	37	24	38
9.9	Linehaul cost factor	37	37	39	25	38

Table 22 Results of experiment 9: Sensitivity analysis – Total routing costs

		Total routing costs				
Exp	**Input parameter**	**50%**	**75%**	**100%**	**125%**	**150%**
9.1	Hub opening cost	3,473	3,585	3,585	3,866	3,642
9.2	PPC duration	3,820	3,746	3,715	3,691	3,625
9.3	HD duration	3,962	3,805	3,723	3,567	3,551
9.4	Fixed vehicle cost	3,452	3,508	3,715	3,932	4,084
9.5	Hourly vehicle cost	2,258	2,919	3,754	4,462	5,219
9.6	Cost per km	3,401	3,714	3,698	3,707	3,913
9.7	Vehicle capacity	3,676	3,721	3,780	3,579	3,866
9.8	Vehicle driving range	3,894	3,646	3,707	3,782	3,547
9.9	Linehaul cost factor	3,750	3,753	3,715	3,785	3,573

Table 23 Results of experiment 9: Sensitivity analysis – Total number of routes

		Total number of routes				
Exp	**Input parameter**	**50%**	**75%**	**100%**	**125%**	**150%**
9.1	Hub opening cost	15	15	15	15	15
9.2	PPC duration	16	16	15	15	15
9.3	HD duration	18	16	15	14	13
9.4	Fixed vehicle cost	15	15	15	15	15
9.5	Hourly vehicle cost	15	15	15	15	15
9.6	Cost per km	15	15	15	15	15
9.7	Vehicle capacity	15	15	15	15	15
9.8	Vehicle driving range	17	16	15	15	15
9.9	Linehaul cost factor	15	15	15	15	15

Table 24 Results of experiment 9: Sensitivity analysis – total route length

		Total route length (km)				
Exp	**Input parameter**	**50%**	**75%**	**100%**	**125%**	**150%**
9.1	Hub opening cost	1,852	2,112	2,112	2,309	2,205
9.2	PPC duration	2,124	2,059	2,057	2,099	1,935
9.3	HD duration	2,391	2,226	2,123	2,107	1,910
9.4	Fixed vehicle cost	2,164	2,252	2,057	2,217	2,217
9.5	Hourly vehicle cost	2,082	2,072	2,125	2,154	2,182
9.6	Cost per km	2,004	2,210	2,056	2,112	2,179
9.7	Vehicle capacity	2,243	2,100	2,217	2,082	2,309
9.8	Vehicle driving range	2,279	2,150	2,139	2,206	2,030
9.9	Linehaul cost factor	2,172	2,170	2,057	2,228	2,041

Table 25 Results of experiment 9: Sensitivity analysis – routing duration

		Total route duration (hr.)				
Exp	**Input parameter**	**50%**	**75%**	**100%**	**125%**	**150%**
9.1	Hub opening cost	104	107	107	117	109
9.2	PPC duration	114	112	113	113	109
9.3	HD duration	118	114	113	108	110
9.4	Fixed vehicle cost	113	109	112	114	114
9.5	Hourly vehicle cost	108	108	114	113	115
9.6	Cost per km	106	115	112	108	114
9.7	Vehicle capacity	110	112	114	107	117
9.8	Vehicle driving range	116	108	112	114	106
9.9	Linehaul cost factor	114	114	112	114	107