A new approach to integrate the routing schedule and the customers' inventory considering shelf life of goods

Author: Benjamin Overkempe

Date: 23-08-2020

Supervisor HAVI: Michiel Degen

First supervisor University of Twente: Martijn Mes

Second supervisor University of Twente: Eduardo Lalla

HAVI

UNIVERSITY OF TWENTE.

Management Summary

HAVI is the logistics provider of various customers in the food industry. McDonald's is HAVI's largest customer and the only customer in scope of this research. HAVI NL manages the inventories of all 250 restaurants. This research focuses on the integration of the disciplines of transport and inventory management. Because McDonald's is a VMI (vendor managed inventory) customer, HAVI can determine the most efficient delivery patterns for them.

Customers' demand has a weekly distribution with increasing volumes towards the weekend. In the weekend, HAVI pays salary supplements, which is included in the labor costs. Furthermore, there are workload peaks in the week as a result of HAVI's previous projects in which the delivery frequency is reduced. These workload peaks result in unnecessary costs.

To determine feasible and efficient delivery patterns, capacities of storages has to be known. HAVI is currently not able to translate the inventory data into restrictions that could be used to optimize the routing schedule. Therefore, the transport planning and inventory management are separated in operations and organizational structure, in the current situation. The goal of this research is to integrate the two disciplines. In this research, we deduced lower bounds for the capacities of the customers' storage locations, expressed in delivery units that can be stored simultaneously. The lower bounds are based on historical data together with the current routing schedule to which we applied forecasted volumes with historical hourly sales distributions. Efficiencies of delivery patterns have to consider geographical optimizations as well as avoiding unnecessary workload at the most expensive shifts with salary supplements (i.e., in the weekend). This research also considers the shelf life agreements for different temperature zones that are made with customers.

This research focuses on (i) reducing the operational costs and (ii) improving the balance of the workload within a week. The main research question is as follows:

How can HAVI save on operational costs by balancing the workload within a week, without violating customers' restrictions including storage capacities at the customers' locations?

In literature, models to optimize routing problems considering inventory policies, are classified as inventory routing problems (IRP). Soysal, Bloemhof-Ruwaard, Haijema and van der Vorst (2015) say that the shelf life restriction is one of the main obstacles to apply IRP models in the food sector. We were not able to identify the literature that optimizes the routing problem and the inventory control simultaneously. The approach suggested in literature is to decompose the problem into subproblems: inventory control and routing. This research is unique, because we determine the delivery patterns and quantities in a single optimization step, which integrates the routing problem and inventory control.

In this research, we both model the tactical as well as the operational routing. We first create a tactical routing schedule, which is a weekly repetitive routing schedule that states which routes are driven each day. This tactical routing schedule is based on deterministic forecasted volumes. After the new tactical routing schedule is built, we model the operational routing, by using historical data from three different weeks. In the operational routing we first determine the delivery quantities given the fixed set of routes from the tactical routing schedule. Second, we improve the operational routing.

In this research, we only use data of customers from the eastern part of the Netherlands, for which we use a maximum of 6 vehicles. we use the tactical routing schedule that is driven from the 6th of July until the 6th of September as the current situation. HAVI uses 22 vehicles for the total distribution network in this period.

To construct the new tactical weekly repetitive routing schedule, we build an algorithm. This algorithm is an improvement heuristic using the current situation as initial solution. We minimize the operational costs, while inventory is included as restrictions. Instead of planning deliveries and determining delivery quantities, we determine the delivery window for each delivery unit. These delivery windows consider storage capacities, shelf life, consumption periods, hourly sales distributions and time windows for the specific customer. An allocation procedure is used to allocate all delivery units to a route. In this procedure we distinguish between focusing on costs first and balancing workload first. We find that the difference between the results is not significant. With our model we have built input for the transportation optimization software, which solves rich vehicle routing problems (RVRP), to optimize the routing schedule.

We experiment with five interventions to improve the tactical routing schedule. Two of the interventions have a significant impact on improving the workload balancing, from which one intervention also significantly contributed to reducing costs. The most effective intervention is to force the RVRP solver to reduce the fleet size.

We find that reducing costs and balancing workload is not a tradeoff. Instead, these are significantly positively correlated. The best tactical routing schedule has the lowest costs and the best workload balancing. In this research, we measure the results in how they contribute to the total network. By changing the distribution network for one fifth of the network, we realized to reduce the fleet size of the total network from 22 to 20 vehicles. The utilization of the trucks is increased from 86% to 89%. Costs are reduced by more than €5000 per week. Additionally, almost €1000 per week can be saved by allowing consecutive deliveries to be within 18 hours of each other. We created a variable to measure the workload balancing. This measure is a penalty stating a normalized value (unitless; score of zero would imply perfectly balanced workload) of the sum of the squared deviations from the required workload balance. By the adjustments made in the eastern part of the Netherlands, in our new tactical routing schedule, we reduced this penalty for the total network from 234,0 to 214,9.

We perform a sensitivity analysis by applying an operational policy that we created, to test how the new tactical routing schedule performs based on actual volumes of three different weeks (volumes deviate from -8,2% to +8,0% in those weeks). We find the new routing schedule being more sensitive to volume deviations. When volumes are 8,0% higher than forecasted, the fleet size of the new tactical routing schedule is equal to the current situation. The operational savings vary from €705 in a week with +8,0% volume to €5406 in a week with -8,2% volume. With the new created routing schedule, more operational adjustments are made in the delivery times and number of routes per shift. We advise HAVI to minimize the number of deliveries that are switched to another shift than they were planned in the tactical routing schedule to minimize the changed delivery times and changes in number of routes per shift.

In the new routing schedule, 18 stockouts occurred in a period of three weeks, compared to 3 stockouts in the current situation. Costs to add extra deliveries are already included in the results mentioned above. Further research should be done on the impact of basing the delivery windows on other volumes. While determining the delivery windows, we consider that the delivery patterns should also be feasible when actual volumes seems to be 110% of the forecasted volumes. To reduce stockouts, this percentage could be increased, at the expense of smaller delivery windows. Furthermore, in this research we consider the shelf life agreements while making the tactical routing schedule. Further research should also be done to incorporate the waste of perishable goods into the objective into the operational IRP with fixed routes.

With our model and research, we helped HAVI to operate more efficient. We have helped HAVI to quantify the capacities of their customers' storage locations. We developed a tactical model, which enables HAVI to take the lead in determining the delivery patterns. This model saves on costs and simultaneously balances the workload more even within a week. We have provided an operational model for HAVI, which guides them to spread the volume over the deliveries, utilizing the existing routes. The VMI partnership with McDonald's is more utilized

We contribute to existing literature with our model and research in several ways. We were not able to identify the literature that jointly optimizes the routing problem and the inventory control. By developing a method in which delivery patterns and quantities are optimized in a single phase, while considering capacity restrictions, we have achieved to integrate the routing problem and inventory control. Simultaneously, we considered shelf life restrictions, which is one of the main obstacles to apply basic IRP models in the food sector. We even managed to include the possibility that shelf life differs per temperature zone. Furthermore, we have achieved to let the delivery quantities be dependent on the delivery times (hour), because we have determined the delivery windows on the level of delivery units. We created this model such that an extended IRP can be solved by an RVRP solver.

Preface

With this master thesis I finish my master study Industrial Engineering and Management, with the specialization in Production and Logistics Management. HAVI has shown to be a great company that has a drive to improve their performances towards their customers.

I want to thank HAVI for letting me do this research for them. I want to thank Michiel Degen for the enthusiasm to share all ins and outs of the business, whether it is in the details or for the greater picture. I'm grateful that due to your passion, I have grown my own passion to strive for the unreachable. The skills you have, to pull me back to see the greater picture again.

I also want to thank my supervisors from the University of Twente. I want to thank Martijn Mes and Eduardo Lalla for the feedback you gave on this work. Martijn, I appreciate your directness, it helped me to be more critical. When sparring about my ideas, you often need half my thoughts to be able to understand what I did not understand yet. Eduardo, I really appreciate your kindness. The feedback you provide is alert. You find missing links and helped me structure parts of this thesis.

I want to thank my wife, Laura Overkempe-Korpershoek, whom without I would not have finished my study yet (maybe, no matter when you read this). Thank you for letting me push harder, and the support you provided. I really am going to enjoy having more spare time again to spend with her, starting with a vacation.

I want to thank all my friends and my family, and a special thanks to my parents who supported me, not only financially, but especially for their love for me. I thank all my dear ones, for all the support during my studies, which already started almost ten years ago. It has been a great time for me to discover my passion, discovered what real friendships are about, and even discovered more of who I am.

I am grateful for the love of God, whom without, this research had no purpose, for all is to be for the glory of Him!

Index

Abbreviations and definitions

3PL: third-party-logistics: Execution of logistic processes by external company.

4PL: fourth-party-logistics: Execution and management of logistic processes by external companies.

CVRP: Capacitated vehicle routing problem: A VRP with a fleet of trucks with given capacities. In this research, we always use a CVRP when discussing a VRP. We use VRP and CVRP interchangeably.

DC: Distribution centre: Warehouse where HAVI stores the goods of its customers before they are being delivered to the customers' locations. Interchangeably we also use depot.

Delivery pattern: A weekly repetitive sequence of days (and times) that represents the delivery days and times for subsequent deliveries for a specific customers' location.

IRP: Inventory routing problem: Routing problem in which inventory levels and possibly costs are included in the objective and/or constraints.

JIT: Just in time: A method within logistics to minimize inventory levels by delivering goods just before the customer needs them.

MLS: McDonald's logistics provider.

Paragon: Transport optimization software used for simulations with deterministic input.

PRP: Periodic routing problem: Routing problem in which a periodic schedule is made, which repeats itself.

Pulling volume: A task executed by a restaurant planner to move a part of the to be delivered volume from a specific delivery to a preceding delivery of the same customer.

Restaurant planner: HAVI employee that represents a set of restaurants from McDonald's and is responsible for making the orders and building a strong relationship with the customers.

RVRP solver: Rich vehicle routing problem solver: Software that is being used to solve VRP problems with a lot of extra restrictions.

Transport planner: HAVI employee that is responsible for planning the day-by-day routes on an operational level.

TSP: Traveling salesman problem. Problem in which the shortest route must be found that visits every point in a given set exactly once.

VMI: Vendor Managed Inventory: A method in which the vendor is responsible for managing the inventory (levels) of its customers.

VRP: Vehicle routing problem: Routing problem in which a given set of orders must be delivered as efficient as possible.

1. Introduction

In Section 1.1 we introduce HAVI first, followed by a description of the main businesses. We continue by stating what transformation HAVI went through the last year. We close the first section with the project description as stated by the project owner within HAVI. In Section 1.2 we formulate the problem from our research perspective, followed by the requirements from HAVI's perspective and the customers' perspective. We continue in Section 1.3 by stating the problem cluster. We explain how the problems relate to each other in Section 1.4. Section 1.5 is used to decide which core problems should be within our scope. In Section 1.6, we decide upon the scope of this research. In Section 1.7 formulate our research goal. In Section 1.8 we conclude with the research approach in which the research questions and research design are discussed.

1.1 HAVI

Worldwide

This research is done at HAVI, named after the founders' wives Harriette and Vivian. HAVI is a worldwide company, executing and managing the supply chains of its customers as a third-party-logistics (3PL) and fourth-party-logistics (4PL) provider. HAVI has grown in the past decades expanding geographically as well as extending to a broader range of services. In 1974, Perlman Rocque first started delivering McDonald's restaurants in the Chicago Metropolitan Area. In 1975, Perseco first started packaging orders from McDonald's. In 1976 HAVI, was formed bundling the services of Perlman Rocque (the two founders) and Perseco to serve McDonald's in the United States.

The Netherlands

In 1986, McDonald's logistics & services (MLS) started to deliver the McDonald's restaurants in the Netherlands from a warehouse managed by MLS. Until 2005, McDonald's was the only customer. MLS was growing fast because of the growth in volume of McDonald's and because it expanded the services with a focus of distressing McDonald's of all activities that were not a core activity for them. Services that MLS has taken over are, a.o., waste and garbage recycling, managing utilities contracts and IT systems. In 2005, British Petroleum (BP) signed as the second customer. Four years later, in 2009, MLS changed its name to HAVI Logistics BV and became part of the worldwide firm HAVI Group. Over the world, HAVI has different customers, but everywhere HAVI operates, it at least delivers McDonald's restaurants.

Next to the core businesses, warehousing and distribution, HAVI offers services in marketing analytics, data analytics, forecasting and packaging solutions to its customers. In the Netherlands, where the research for this master thesis takes place, HAVI has two warehouses, in Amersfoort and Barendrecht. From those locations, mostly food products are delivered to ten customers with a total of around six hundred locations.

In this research, we focus only on the distribution from the warehouse in Amersfoort. The warehouse in Amersfoort delivers to all McDonald's restaurants, which is the most interesting customer for this research. The reason for this is that HAVI has the most influence on this customer because to McDonald's, HAVI is a 4PL provider (i.e., HAVI is taking over a large part of the supply chain of McDonald's including next to transport and warehousing also inventory management and forecasting) in a vendor manages inventory (VMI) partnership. From Amersfoort, HAVI delivers to 250 locations of McDonald's and 120 locations of BP. The goods for the other customers are stored in the warehouse in Barendrecht. Since 2020, HAVI also delivers some of these customers' locations from the warehouse in Amersfoort after the products are cross-docked from the Warehouse in Barendrecht to Amersfoort. The total distribution from Amersfoort is executed with 22 own trucks.

HAVI's structure change in 2019

In the year 2019, HAVI went through an organizational structure change. Instead of a separation by discipline (i.e., transport and warehousing were separate pillars), the new structure separates the company on decision making level. Therefore, the disciplines are taken together on strategic, tactical and operational level. As stated by Potter, Towill and Disney (2007), transportation is often optimized on its own, within the constraints enforced by the supply chain. Although, there has been an increased recognition that transport needs to be integrated into the supply chain. As the structure change within HAVI is still recent, the first steps are taken to align the disciplines on the strategic level. The next step is to integrate the disciplines on the tactical level.

Project Description

HAVI has given the assignment to design a model or procedure that is able to guide them to balance the workload for transportation and warehouse within a week. The peaks experienced result in unnecessary costs because a larger fleet is needed and inefficiencies in the warehouse occur. When HAVI is in the lead of the delivery patterns of their customers, which enables them to influence the balancing of workload, large potential benefits in terms of operational costs can be achieved.

1.2 Problem statement

HAVI wants to integrate the disciplines of transport and inventory management to gain more benefits from their VMI partnership with McDonald's. More specific, HAVI wants to include the knowledge and insights of inventory management at their customers' locations into the transport optimization. Waller, Johnson and Davis (1999) state that in a VMI partnership, the vendor makes the main inventory replenishment decisions. The vendor is responsible for the buyer's inventory levels. The benefits of VMI are categorized in reduced costs and improved service. Costs are reduced at the inventory of the supplier as well as the locations of the consuming organization. Also, transportation costs are reduced.

So far, HAVI has not been able to translate the implications of the inventory limitations to constraints and restrictions within the routing process. The customers of HAVI expect HAVI to become more efficient every year. In 2019, a project with the headquarters of the customers took place to create a better routing schedule, which is summarized as follows:

- Simulation studies showed potential annual savings of €500.000.
- Delivery patterns chosen by HAVI instead of the customer.
- Simplistic by the assumptions that were made.
- Either the potential savings or the customers' wishes were fulfilled, not both.

The main reason HAVI is not able to create a routing schedule that fulfils customers' wishes and results in comparable annual savings as in the simulation studies, is that they do not know how to include the data and knowledge about the inventories of their customers, into the transport optimization model.

Integrating transport and inventory management

HAVI is a 3PL provider for all its customers and a 4PL provider for its largest customer, McDonald's. From the warehouse in Amersfoort, McDonald's accounts for 84% of volume transported (Measured in delivered roll containers in 2019). HAVI as a vendor, manages the inventory of every restaurant of McDonald's in the Netherlands. HAVI should be able to integrate the transportation and inventory management for its largest customer. Although HAVI manages the transport and the inventory of every McDonald's restaurant, these two disciplines are not managed together, but separately. Potter et al. (2007) describe the potential benefits of integrating transport into the supply chain:

- Improved customer service levels.
- Lower transport costs.
- Improved vehicle utilization.

HAVI's requirements

HAVI's requirements for this research, is split in transportation requirements, warehouse requirements and management requirements. Many of the requirements for this model are the same requirements that apply to making the tactical routing schedule. Therefore, we will state them briefly.

Transport requirements

- Time windows are used (see Appendix I).
- Routes should not be optimized on duty time (i.e., departure from depot until arrival at depot, increased with one and a half hour for (un)loading and administration) alone, but also include clustering based on geography.
- A buffer time must be included between the arrival of the day shift and departure of the evening shift, such that delays during the day do not have an impact on the evening deliveries
- Mitigate workload peaks (more elaborated in Section 1.4, see Figure 3).

The routes should be (i) cost efficient and (ii) be distributed such that when volume increases or decreases, the transport planners can adjust the routes without too many changes. Besides these standard requirements, the requirements specified for this research are the following:

- The demand should be balanced such that storage capacity restrictions are met.
- There should be an operational policy that guides the restaurant planners in which volume to pull to earlier deliveries.

Warehouse requirements

- The workload peaks are mitigated.

The most important goal for the warehouse is to mitigate the peaks, because volumes that are higher and lower, both have their reasons why they bring inefficiencies, explained further in Section 1.4.

Management requirements

- The workload peaks are mitigated.
- HAVI is in the lead of determining delivery patterns instead of the customer.
- Save on operational costs.
- Reduce pollution.

From a management perspective, there are multiple reasons to mitigate the peaks of workload. With a more even workload, the inefficiencies in the warehouse decrease, and a smaller fleet size is needed (which is a large cost component) to execute the distribution. To lower the operational costs, workload should be minimized on the days on which salary supplements apply. Instead of letting the customers determine the delivery patterns, HAVI should determine the delivery patterns that results in the most efficient operations. This way, HAVI can save on operational costs.

Customers' requirements

- Time windows are used. These time windows exclude the times that deliveries are interrupting in such a way that either the core operations or the customers are hindered.
- The delivered quantity should always fit in the storage.
- The delivered goods should always have a shelf life with a minimum of what is agreed upon.
- When a stockout is likely to happen, an extra delivery has to be planned to replenish the stock.

Furthermore, from the customers perspective, the less changed delivery times, the better. HAVI measures changed delivery times, by counting the number of deliveries that are planned in the operational route more than 30 minutes earlier or later than it was planned in the tactical routing schedule.

1.3 Problem cluster

To get an overview of all the related problems, a problem cluster is created, which is shown in Figure 1. The red problems are the end problems (i.e., the effects). The yellow problems are the problems that we cannot influence or are obviously out of scope. The blue problems are the problems for which we do not have a further cause upstream in the cluster, or they only have causes that we cannot influence on or are out of scope. Therefore, the blue problems are the core problems, when being solved will result in a chain reaction to the problems further down the stream. The white problems are the rest of the problems. They have a cause that can be treated and they also cause another problem. With green rectangles we highlight the problems that we are going to use to measure the impact of the solution of this research and which are the action problems. The end problems are quite trivial; almost every company in the same branch wants to improve on them. It is the causal problems that differ per company. The end problems are divided into categories based on the main problem owner:

HAVI

- Unnecessary high costs.
- The mistrust from the transport department towards the restaurant planners to balance volume amongst the days within a week.
- The restaurant planners being busy with non-core tasks.

Customer

- The restaurant planners being busy with non-core tasks.
- Changed delivery times.
- More pollution. (Both McDonald's as BP have impact on the environment high on the agenda)

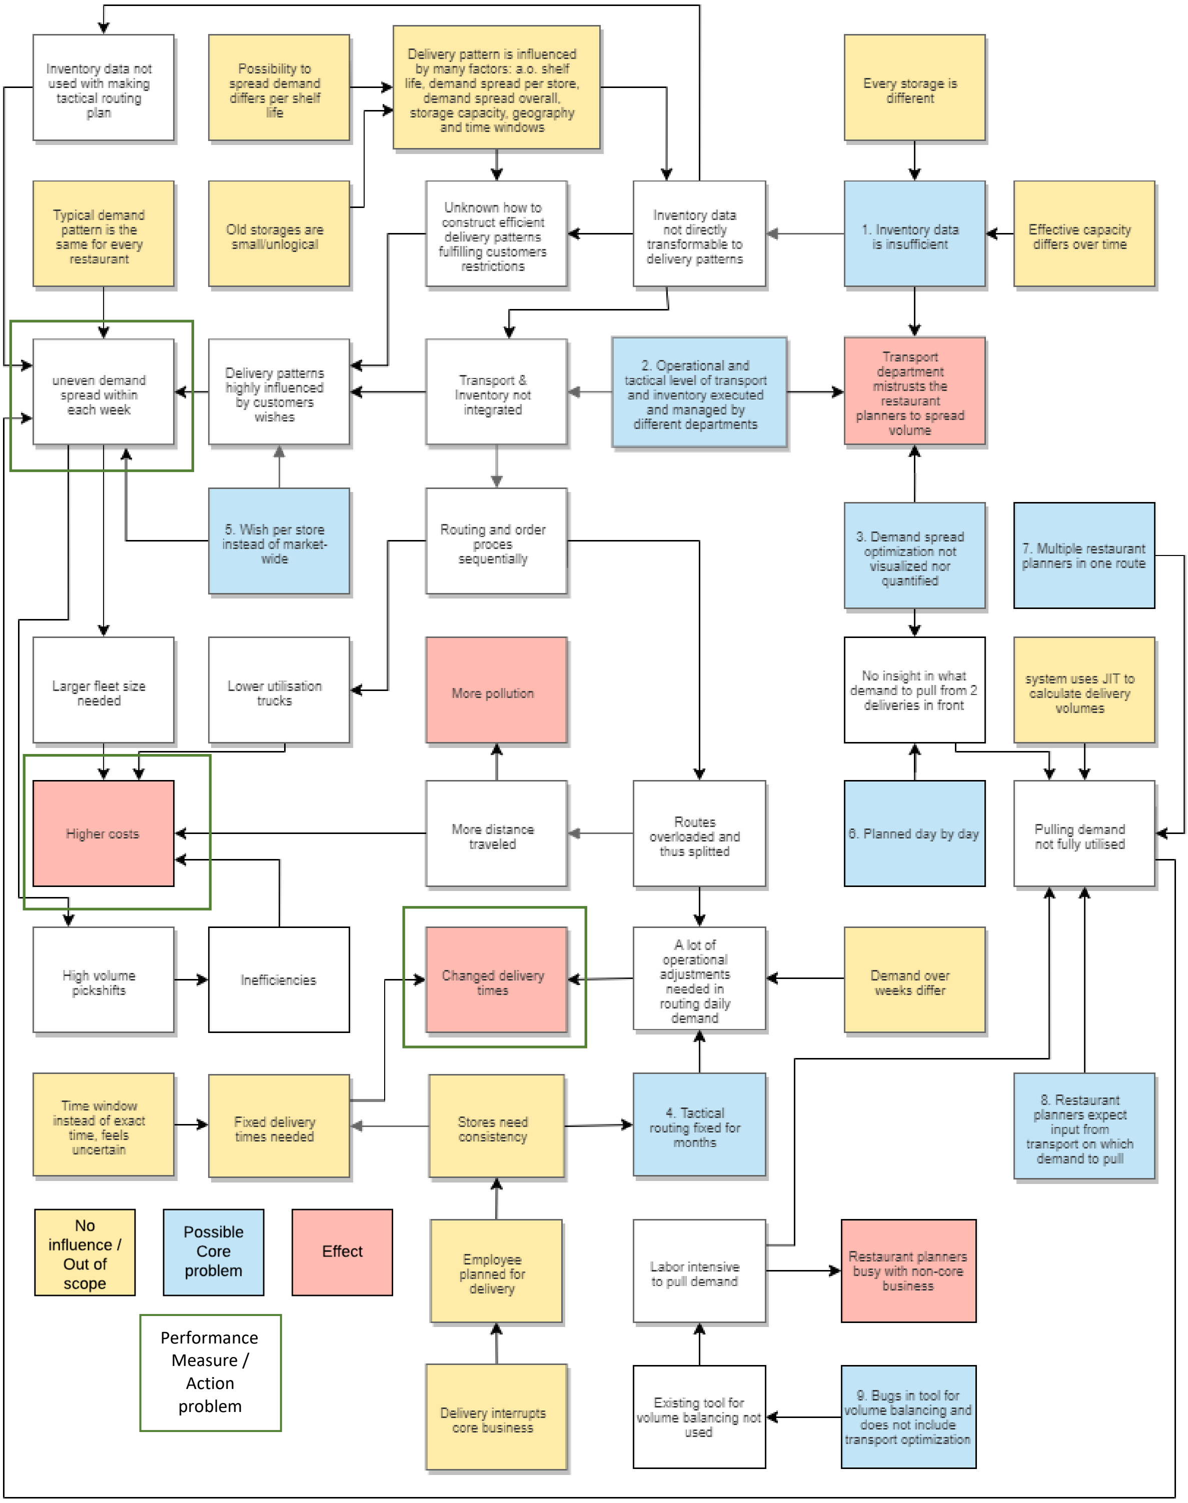

Figure 1: Problem cluster

The main problems, as the project owner defines it, are high operational costs and uneven workload balance. When moving more downstream from 'uneven workload', the end problem is again higher costs.

1.4 Core problems

Nine core problems are detected. The nine core problems are as follows:

1. The available inventory data is not sufficient to determine the actual storage capacity and construct feasible and efficient delivery patterns.
2. On tactical and operational level, the transport and customers' inventory are executed and managed by different departments.
3. The optimization of demand balancing is not visualized, nor quantified.
4. The tactical routing schedule is fixed for a period of months.
5. Wishes per restaurant are highly influencing the current delivery patterns, although HAVI should be in the lead to create more efficiencies.
6. The transport planners and the restaurant planners act day-by-day.
7. Having restaurants of multiple restaurant planners within one route, makes it cumbersome to discuss which volume (i.e., part of an order) must be pulled to preceding deliveries.
8. The restaurant planners expect input from the transport department on which volume to pull, but do not get that input.
9. The tool that is developed to balance volume has bugs and does not consider transport optimization.

We must choose one or multiple of these core problems, which we want to solve to have an impact on the end problem(s). The end problems as stated by the project owner are directly or indirectly higher operational costs. The problem cluster in Figure 1 shows a clear overview of all the problems. To focus on one or multiple core problems, it is necessary to know the significance of the problems and some more details of their context. In the coming paragraphs, we discuss the problems and explain how they relate to each other. We group the problems based on the subject they are related to. We first discuss the problems directly relating to higher costs, because that is related most closely to the overall goal. Then we discuss the problems from an organizational view (i.e., how the management structure contributes to the problem). Following, a discussion of the problems relating to an uneven workload. Thereafter, we discuss the lack of insight in balancing volume, followed by why volume is not balanced to maximum potential, first by the system, and second by human actions. Next, we discuss the problems that HAVI has faced at earlier attempts to solve this problem. We end with discussing the causes of failing to include the information about the inventories at the customers' locations (i.e., why it is so complex to create feasible and efficient delivery patterns).

Unnecessary high costs

HAVI wants to lower the transport costs without having a negative effect on the customer service levels. Song and Savelsbergh (2007) state that the best measure to compare different solution approaches for instances dealing with routing problems considering the inventory levels at the customers, is the volume per kilometer. This measure can be improved by either:

- Increase volume per route (higher utilization).
- Decrease distance traveled per route (geographical efficiency).

Therefore, utilization is not the only cost driver, but it is often a good indicator of the relative transportation costs given the volume transported. To be more specific, the transportation costs are split in the two main factors, namely: the amount of fuel used, and the time needed for the driver to execute the route. These two are highly positively related to each other. The distance traveled, which has the largest influence on the fuel used as well as the duty time of the truck driver, therefore are good indicators for the transportation costs. Another major cost factor for HAVI's operations is the costs relating to owning trucks (i.e., lease contracts, insurance a.o.). These costs are relevant for this research because mitigating workload is one of the main requirements and the fleet size is dependent on the shifts with the largest volume.

Organizational view

Looking at integrating the transport and inventory management, we must determine who is responsible for which activities and how these disciplines are managed. The transport planners and restaurant planners are physically separated, because they have their own office. Looking at the organizational chart (see Figure 2), the restaurant planners and transport planners are related to each other via the Managing director NL. The departments have daily contact, to solve issues. However, to deal with conflicting goals, agreements on higher levels in the organization structure must be made. The managing director cannot be involved in those matters and the two disciplines do not cooperate intensively when it comes to long-term solutions. This results in two disciplines that blame each other for acting in their own favors. To align the departments, a role is needed lower in the organization that is responsible for both departments, or they have to work together in projects with stakeholders who make agreements and discuss progress in meetings.

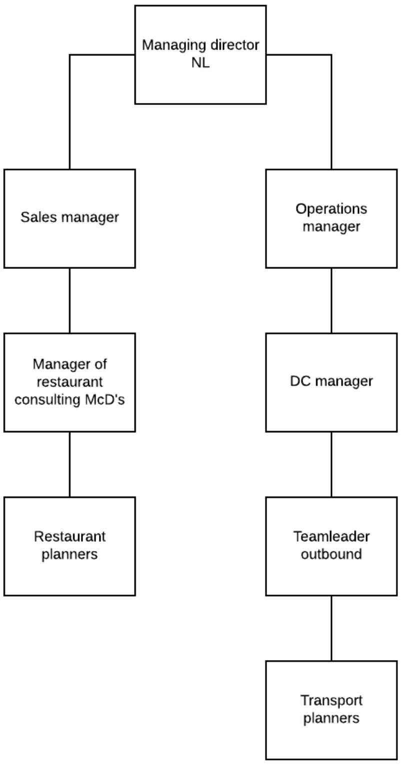

Figure 2: Organizational chart

Uneven workload

One of the main goals for HAVI in this research, is to flatten the workload within a week. An uneven balance has many disadvantages, from which the biggest impact on costs is that fleet size is based on maximum number of trucks needed during a shift. In Figure 3 we show the balance of the consumption volume of all McDonald's restaurants. We also show that the volume that is being delivered from Monday to Friday is higher than what is consumed. Saturday and Sunday, the delivered volume is smaller than the consumed volume.

In Figure 3 we show what intuitively would be the ideal balance considering an even workload throughout the week with decreasing volume in the weekend because of the salary supplements (i.e., 150% pay on Saturday and 200% pay on Sunday).

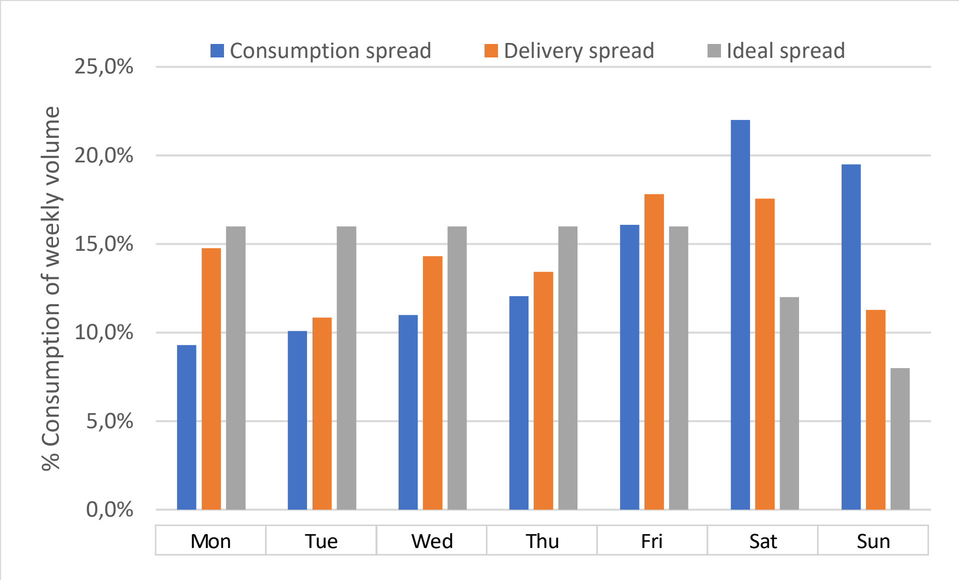

Figure 3: Volume balancing McDonald's across the week

We take volume of goods (measured in delivery units) as indicator for workload. For transportation this is a good indicator, because the characteristics of HAVI and its customers are such (routes with few and relatively large stops) that the capacity of the truck is the main limitation instead of the maximum duty time. When volume increases, the number of routes needed increases approximately linearly. In the warehouse, workload also increases practically linearly with volume. Therefore, volume is a good indicator for workload, and balancing workload and balancing volume are used interchangeably in our case. The main implications of the different balances are:

- Fleet size can be decreased.
- Storage spaces at the customers' locations are being filled more in the beginning of the week and emptied more in the end of the week.

To quantify the impact on the customers' storage, we introduce the term Center of Gravity (CoG). The CoG for deliveries is the weighted average delivery moment in the week. The center of gravity in this context is defined to be

Equation 1

$$CoG = \frac{1}{D} * \sum_{i=1}^{7} i * n_i$$

where i is the day number of the week in which Monday equals 1 and Sunday equals 7. The number of delivery units that are delivered on day i is denoted by n_i and the total demand in the week is denoted by D. The choice of setting Monday equal to 1 and Sunday to 7, is not arbitrary. On Monday, consumption is lowest, and increases throughout the week. Determining the CoG for deliveries, we find a CoG of 4.06. Determining the CoG for consumption we find a CoG of 4.59. This means that consumption takes place more later in the week compared with the deliveries. For the customer it is favorable that the CoG of consumption and the CoG of delivery are closer to each other, such that they do not store the goods longer than needed. In Appendix II, we elaborate more on the CoG, and how it is influenced by choosing other delivery patterns.

Besides the extra costs related to transport and the customers' storage, the peaks in the workload also results in extra costs in the warehouse:

- Limited space in the chilled and frozen area causes order pickers to block each other's paths.
- In the dry area the workload is limited by the capacity of the forklift trucker to replenish the pick locations with new materials from the bulk storage.
- The pickers more often see empty pick locations and must wait or must pick later to complete the order.
- With low volumes still one FTE forklift trucker is needed, but his utilization is much lower.

Not balancing to full potential

After introducing the importance of balancing the workload, we now discuss the various reasons why the workload is not being balanced to full potential:

- The system of HAVI that proposes the orders uses Just-In-Time (JIT) delivery.
- The actions that are (not) taken.
- There is no insight in which part of demand can be delivered earlier, and how much capacity per storage space is available to deliver more than strictly needed.
- Volume balancing tool does not integrate transport.
- McDonald's is the only VMI customer for which the order quantities are determined by HAVI.
- Lack of knowledge when volume deviates from forecast.
- Day by day planning instead of looking ahead.
- Complexity of delivery patterns.

The system

The system of HAVI that proposes the orders uses JIT delivery. This means that the delivery quantity is totally dependent on the amount of forecasted consumption between two consecutive delivery moments. As shown in Figure 3, the balancing of consumption differs from the ideal balance, and thus delivery quantities should not be determined solely by the consumption amount. In Appendix III we show what the impact is of the utilization of routes given that the system uses JIT order calculations.

The actions

The restaurant planners are able to adjust the orders that are proposed by the system. They can pull volume to an earlier delivery than the system has planned. The effect is shown in Figure 3 because the delivered balance is more flattened relative to the consumption balance. Because the task of pulling volume to earlier deliveries is time consuming, the restaurant planners do not do this when it is not directly necessary. Pulling demand is time consuming because (i) an order contains about two hundred different products. For every product must be determined to pull the volume to an earlier delivery and (ii) the restaurants within a route can be represented by different restaurant planners, thus they must consult with each other which volume to pull to earlier deliveries. Furthermore, there is no real-time insight in the result of pulling volume to earlier deliveries.

Insight

HAVI has no insight in what the potential is of flattening the delivery balance further. Figure 3 shows how the ideal balance would look like, but it is not known whether this can be achieved or how close this can be matched.

Volume balancing tool

The supply chain department has designed a tool that is called 'Volume Balancing', which could be used by the restaurant planners. The tool 'Volume Balancing' is made by the IT team at the headquarters of HAVI in Germany. The orders proposed by the tool are not based only on JIT, but also incorporates volume balancing throughout the week. The tool is tested but proposes orders that were not as they should have been and therefore is not used anymore. Besides, this tool is an example of how transport and ordering are two separated disciplines within HAVI. Because transport optimization is not incorporated in this tool. The volume is balanced only based on volume per customer, independent on how the routes are organized. Therefore, this tool is not used in this research.

McDonald's as only VMI customer

The delivered goods for McDonald's and the other customers are delivered in the same routes. Because the other customers determine the delivery quantities themselves, HAVI is not able to fully utilize the capacity of the truck, because the rational used by the customers, is not easy to understand.

Deviation from forecast

When actual volumes are higher than forecast, the restaurant planners do not know whether to pull volume to earlier deliveries or not. Pulling volume can result in an extra route that must be created. Not pulling can result in problems later on in the week. When volumes are lower than forecasted, the restaurant planners do not have the urge to pull volume to earlier deliveries; the potential of saving a route can be achieved be either pulling or not pulling demand, dependent on the rest of the routes.

Day by day planning

When deciding on whether to pull volume to earlier deliveries, the restaurant planners only look at the deliveries of tomorrow and the consecutive deliveries. However, for most goods stored in the frozen and dry area, the shelf life is longer and thus potentially more volume can be balanced amongst deliveries when looking further ahead. Sometimes it is even necessary as can be seen in the following example and shown in Table 1. The capacity of the trucks is 60 delivery units. Before any demand is pulled to earlier deliveries, the routes have a planned utilization as can be seen in the initial situation. The Friday route is planned above capacity. Therefore, demand must be pulled to an earlier delivery. When only looking one delivery ahead, the result is the second line, and the Friday route is still planned above capacity. When looking multiple deliveries ahead, the restaurant planner can pull volume from the Friday delivery to the Monday delivery.

Table 1: Horizon of the restaurant planners, looking ahead from Sunday

Situation	Route Monday	Route Wednesday	Route Friday	Action required
Initial situation	45	55	75	Pull volume
Acting day by day	45	60	70	Split Friday route
Looking ahead	55	60	60	None

Complexity of delivery patterns

The delivery patterns for McDonald's are determined by the restaurant planners. The other customers determine their delivery patterns themselves. The delivery patterns of the other customers are out of scope for this research and are assumed to be fixed, equal to the current situation. Most of the decisions are based on intuition and preferences from the customers, instead of data analysis. The objectives for HAVI and the restaurants are different.

HAVI

- Flatten workload.
- Geographical efficient routes.

McDonald's Restaurants

- Receive deliveries at convenient moments.
- Minimize the storage utilization fluctuations.
- Minimize number of interruptions in their core business.
- Not too large gaps between two deliveries.
- Storage capacity restrictions are met.

Proposals for new delivery patterns are often refused by restaurants with arguments based on intuition and feelings. Because the restrictions are not stated in writing, and the data HAVI uses is not sufficient, HAVI is not able to verify whether the arguments of the restaurants are legit. Delivery patterns must be created feasible and efficient, both are influenced by many factors. We sum up the most important factors.

Feasible delivery patterns:

- Capacity per temperature storage *(hard restriction).*
- Time windows *(hard restriction; see* Appendix I*).*
- Demand balancing curve *(input).*

Efficient delivery patterns:

- Balance of volume over all customers within a week *(objective: balancing).*
- Positive correlation with neighboring delivery locations *(objective: costs).*

In this research, we must create feasible delivery patterns. The objective is to create efficient delivery patterns. The objectives of HAVI are considered in combination with the storage restrictions of the customers' locations.

In the simulation studies that are mentioned in Section 1.2, HAVI used the following data to determine the delivery patterns:

- Surface of storage location per temperature zone per restaurant.
- Hourly demand forecast distribution per restaurant.

Because this data turned out to be insufficient, the delivery patterns are not yet determined based on data, but highly influenced by customers' wishes.

1.5 Decision on Core Problem

We started Section 1.4 with stating the nine core problems. After the elaboration of the problems, we understand the context and significance of the different problems and how they relate to each other. In our opinion it all comes together in the fact that transport optimization and the ordering process are dealt with separately. A volume balancing tool (Core problem 9) would work for making orders or planning multiple days ahead (6) can help to improve the processes, but it will not be effective when the two disciplines are still dealt with separately. A visualization or quantification of how workload can be balanced (3) is helpful to get insights, but on its own it will not change the way of working. Because we focus on the combination between the tactical and operational level, the organizational structure (2) is out of scope for this research. The tactical routing is being fixed for a period of months (4). Changing this aspect, also requires major organizational changes, which is out of scope for this research. The core problems we focus on are:

1. The available inventory data is not sufficient to determine the actual storage capacity and construct feasible and efficient delivery patterns.
5. Wishes per restaurant are highly influencing the current delivery patterns, although HAVI should be in the lead to create more efficiencies.
7. Having restaurants of multiple restaurant planners within one route, makes it cumbersome to discuss which volume (i.e., part of an order) must be pulled to preceding deliveries.
8. The restaurant planners expect input from the transport department on which volume to pull, but do not get that input.

The first and fifth core problem relate to the tactical level and are related to each other. Both imply that the way in which delivery patterns are constructed is not good. The model we build should construct the delivery patterns while meeting the storage capacity restrictions. The seventh and eight core problem relate to the operational level. When making the operational policy we must guide the restaurant planners how to balance the volume.

1.6 Scope

The primary scope of this research is the process of making a tactical routing schedule, which is a weekly repetitive schedule determining which routes to be driven each day. This weekly routing schedule serves as a basis for the operational routes every week within the planning horizon. In this research, the planning horizon is the period from the 6th of July until the 6th of September. The operational policy is used to check the results with historical data.

McDonald’s is the only customer in scope for this research because HAVI is a 4PL provider to them. We must consider the other customers, because they are distributed in the same set of routes, but the other customers are treated the same way as in the current situation. For them, the delivery patterns and quantities are being fixed and equal to the current situation.

We selected 79 customers that are located in the east of the Netherlands, which is approximately one fifth of the total network. First, we selected all customers in the city of Enschede. Then we selected all customers who are in the current situation at least once routed together with one of those selected customers. We then again selected all customers who are at least once routed together with one of the selected customers. After these two iterations, 50 McDonald’s locations and 29 locations of other customers are selected. The rest of the locations are out of scope. When looking at the results of the experiments we do score the routing schedule based on their contribution to the total network.

We do not optimize the algorithm that is currently used by HAVI to solve the rich vehicle routing problem (RVRP). We define the input to escape from local optima, and input extra restrictions to include the objective of balancing workload.

1.7 Research goal

HAVI wants to balance the workload within a week. To do so, they must improve their tactical transport planning as well as the operational policies used in the daily planning. The tactical routing schedule should be such that workload is more balanced while meeting customers' storage capacities restrictions. The next section states the research questions that need to be answered to know how HAVI can utilize the sales forecasting data, to minimize the total relevant costs, while meeting all the restrictions such as inventory levels at the stores. The main research question is stated below.

How can HAVI save on operational costs by balancing the workload within a week, without violating customers' restrictions including storage capacities at the customers' locations?

1.8 Research approach

To answer the main research question, we formulated eight research questions and twelve sub questions.

Research questions

1. How does the current situation of HAVI look like?
    a. How does HAVI execute the transport optimization?
    b. How does the order process of HAVI look like?
2. How is volume or workload balancing dealt with in the models in literature (i.e., PRPs and IRPs)?
    a. Which solution approaches are used to solve the models?
    b. How do these approaches deal with volume or workload balancing?
3. How can the requirements be translated into restrictions or inputs that can be used in routing optimization software?
4. How can workload balancing be applied in the tactical routing process of HAVI?
    a. What is the objective of HAVI that should be used in this research?
    b. What input is required to set up the model that will be implemented into the processes of HAVI?
5. How to build a model to optimize the transport and inventory in an integrated way on the tactical level?
    a. How does the model differentiate relative to the known models from literature?
    b. Which steps are used in the model?
6. How can the tactical model be tested using experiments?
    a. Which scenarios must be tested?
    b. How to compare the results of the experiments?
7. What operational policy should be followed after implementing the new tactical routing schedule?
    a. What operational policy should the restaurant planners follow?
    b. What operational policy should the transport planners follow?
8. What performance can be achieved by applying the new tactical model and operational policies into practice at HAVI?

Research design

Before doing any research, we have to understand the business, therefore, the first research question is stated to understand the current situation. The current situation is split in transport optimization and the ordering process, because these two aspects together determine how workload in balanced for HAVI. In our research we want to build upon existing literature, therefore, the second research question is to perform a literature study. In this study we focus on the workload balancing in different approaches used. The first of four core problem that we focus on is about the data from customers' storages that is not translatable for the purpose of routing optimization. The third research question is to facilitate this translation. The second core problem in our focus, is about taking the lead in determining the delivery patterns to create efficiencies. These efficiencies have to result in a workload that is more balanced. Therefore, the fourth research question is about achieving the balancing of workload within the tactical routing schedule. We then continue with the fifth and sixth research in which we build and test the model that we are going to use for the tactical routing schedule. The third and fourth core problem in our focus are about deciding on an operational level, which volume to pull to earlier deliveries. The seventh research is therefore to define the operational policies that should be applied to gain the most efficiencies. The eighth research question is to combine the tactical routing schedule with the operational policy to know what results can be achieved when both are in place.

Research approach

RQ1. This question will be answered in Chapter 2 by describing the current situation.

RQ2. This question will be answered in Chapter 3, by carrying out a literature review.

RQ3. This question will be answered by translating the verbal requirements into mathematical formulated restrictions. This question will be answered in Chapter 4.

RQ4. Part a will be answered through a combination of input from the management of HAVI and the findings from the literature study. Part b is answered building further on the results of RQ3. Based on the requirements we determine which input is needed to meet those requirements. This question will be answered in Chapter 4.

RQ5. Based on the results of RQ3 and RQ4, we develop a model. We first determine the delivery windows for the delivery units. We then develop a method to deduce the capacity of the storages based on two criteria. Furthermore, we adjust the input data such that the current situation is a feasible situation. We build an algorithm which optimizes the routes while the inventory restrictions are considered. This question will be answered in Chapter 4.

RQ6. Paragon, the routing software HAVI uses, will be used to execute the experiments. Based on the results of RQ3, we develop spreadsheets models, which will be used to generate input for Paragon. Based on the specific requirements and objective that results from RQ2 and RQ3 we determine the set of experiments that will be executed. At least a re-optimization of the current situation within Paragon must be included in the experiments to make well founded conclusions. To answer part b, we use an efficient frontier to tradeoff costs and balancing workload. Part a will be answered in Chapter 4. Part b will be answered in Chapter 5.

RQ7. To answer this question, we need to determine which volume is pulled to earlier deliveries. We define a set of rules, when followed determine which volume is pulled to earlier deliveries. These policies are used to examine the operational results of the routing schedules that lie on the efficient frontier. The policy c applied when actual order and consumption volumes are known. We will use historical data to test how the new created tactical routing schedule would have performed. This question will be answered in Chapter 4.

RQ8. To answer this question, we need to combine the results of the tactical routing schedule of Chapter 5 and apply the operational policy we constructed in Chapter 4. We perform a sensitivity analysis with historical data as an example of what the volumes in a certain period will be. Although we have all information available in advance, we need to apply the operational policy that results from RQ7 and act only on data that was available at that time. Furthermore, we apply the operational policy on the base case scenario. This question will be answered in Chapter 5.

2. Current Situation Analysis

In this chapter, we answer research question 1: '*How does the current situation look like?*'. To understand the problem in more detail, in Section 2.1 we give some figures and numbers of HAVI's network in the Netherlands. Followed by a description of the current transport and its optimization software and inventory management processes of HAVI in Section 2.2 and 2.3 respectively. In Section 2.4 we discuss how the delivery patterns for McDonald's are determined in the current situation. In this chapter we discuss the current situation of the transport and ordering/inventory management activities and the integration between those disciplines.

2.1 Network and Figures

In Figure 4, we show all customers' locations. The yellow triangles represent the DCs. The black circles represent the customers from which the goods are stored in the DC in Amersfoort (Centre of the Netherlands). The pink circles represent the customers from which the goods are stored in the DC in Barendrecht (South-West of the Netherlands). Since the Covid-19 crisis some of the customers, which

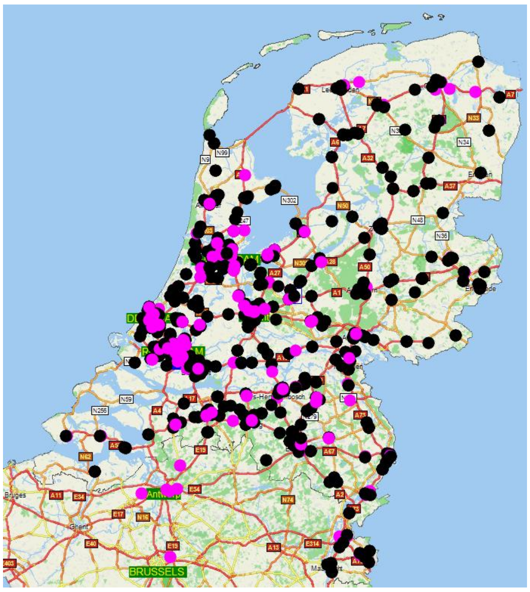

Figure 4: Customers' locations

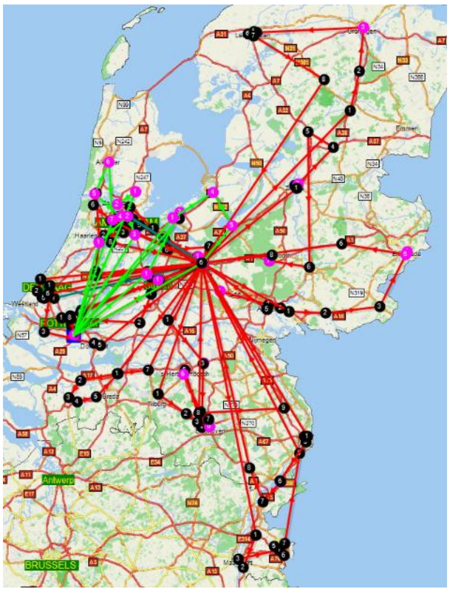

Figure 5: Routes of one shift

goods are stored in Barendrecht, are delivered from Amersfoort after they have been transported from Barendrecht to Amersfoort.

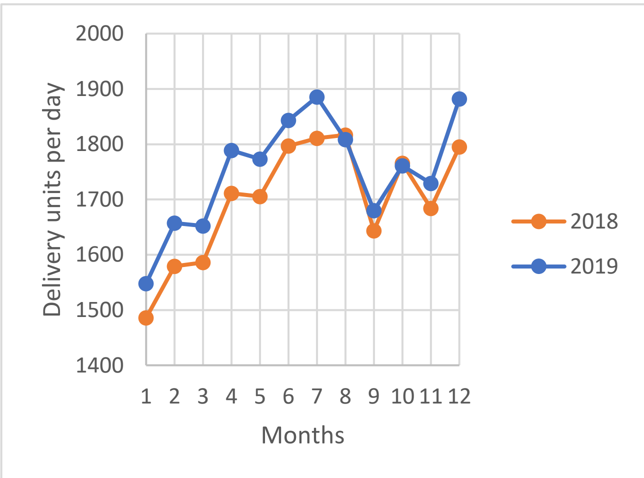

Figure 6: Volume seasonality's

Figure 5 shows an example of all the routes in one shift. The red lines represent the routes departed from Amersfoort and the green lines the ones departed from Barendrecht. From Amersfoort, routes in the dayshift depart on average around 5 :30 am and arrive at 2:35 pm. In the evening shift, routes depart on average around 4:20 pm, and arrive 12:05 am. On average seven customers are in a route from Amersfoort, for Barendrecht the average number of customers is around 4. This difference is due to the capacity of the truck. The distribution in Amersfoort is executed by own drivers, in Barendrecht the distribution is outsourced. Figure 6 shows the seasonality of the volume throughout the year. The summer months are the busiest months. Another remarkable month is December, which is a special month with a lot of discounts, which increases volume. Especially the first months of the year, volume is low. Due to Corona, the utilization rate during 2020 is exceptional. The average utilization in 2018 and 2019 was between 84% and 85%. Besides the volume changes within a year, we also cope with volume changes within a week, which we already discussed in Chapter 1 (see Figure 3).

2.2 Transport optimization

The transport department of HAVI is separated in employees who are responsible for the daily operation (i.e., operational planning), and employees who are responsible for the tactical and strategical time horizons. The future is split into two horizons. The first horizon is a period of weeks to months. This period is the tactical planning period. The second horizon is a period of months to years. This period is called the strategic horizon. For this horizon simulation studies are executed to get answers to what-if scenarios for example.

Tactical Planning

HAVI uses a standard routing schedule, consisting of one week, which forms the basis for a certain period (i.e., a periodic schedule). This schedule is repeated every week, until the standard routing schedule is changed again. The standard routing schedule is adjusted when volumes are expected to change, about four to six times a year. For example, after the summer, both customers decrease in volume, because they are both positively correlated to the weather. The procedure of adjusting delivery frequencies per store, merging and splitting routes, which are based on a seasonal change, is the tactical planning. We made an overview of the process of developing a new tactical routing schedule and present it in Figure 7. The process of creating a new tactical routing schedule starts with gathering input from various stakeholders represented by the green cells in Figure 7. Thereafter an engineer starts with making an initial solution for a new tactical routing schedule. When specific input results in severe impact, the engineer consults with the input giver. After the initial solution is build, the engineer and a transport planner adjust the schedule together after which the plan is communicated internally. The various stakeholders give feedback, and this will be processed. After communicating the plan internally for the second time, the customers are also being informed about the new routing schedule. Feedback from the customers is also being processed after which the changes are communicated to the customer again.

Operational Planning

Next to the seasonality, the volume is affected by much more factors, which are not all predictable. The demand of a set of customers, that form a route, can be more than the capacity of a truck in some weeks. Then, a customer must be transferred to another route, a new route must be created, or a customer must be delivered by an external carrier. Other weeks, the demand of the same set of customers is less than the capacity of a truck. If multiple routes in the same region have overcapacity, these routes must be merged to avoid inefficiencies. These daily adjustments, after the orders of the customers have been approved, is the operational planning. The importance of an accurate tactical planning is the minimization of needed adjustments in the operational planning.

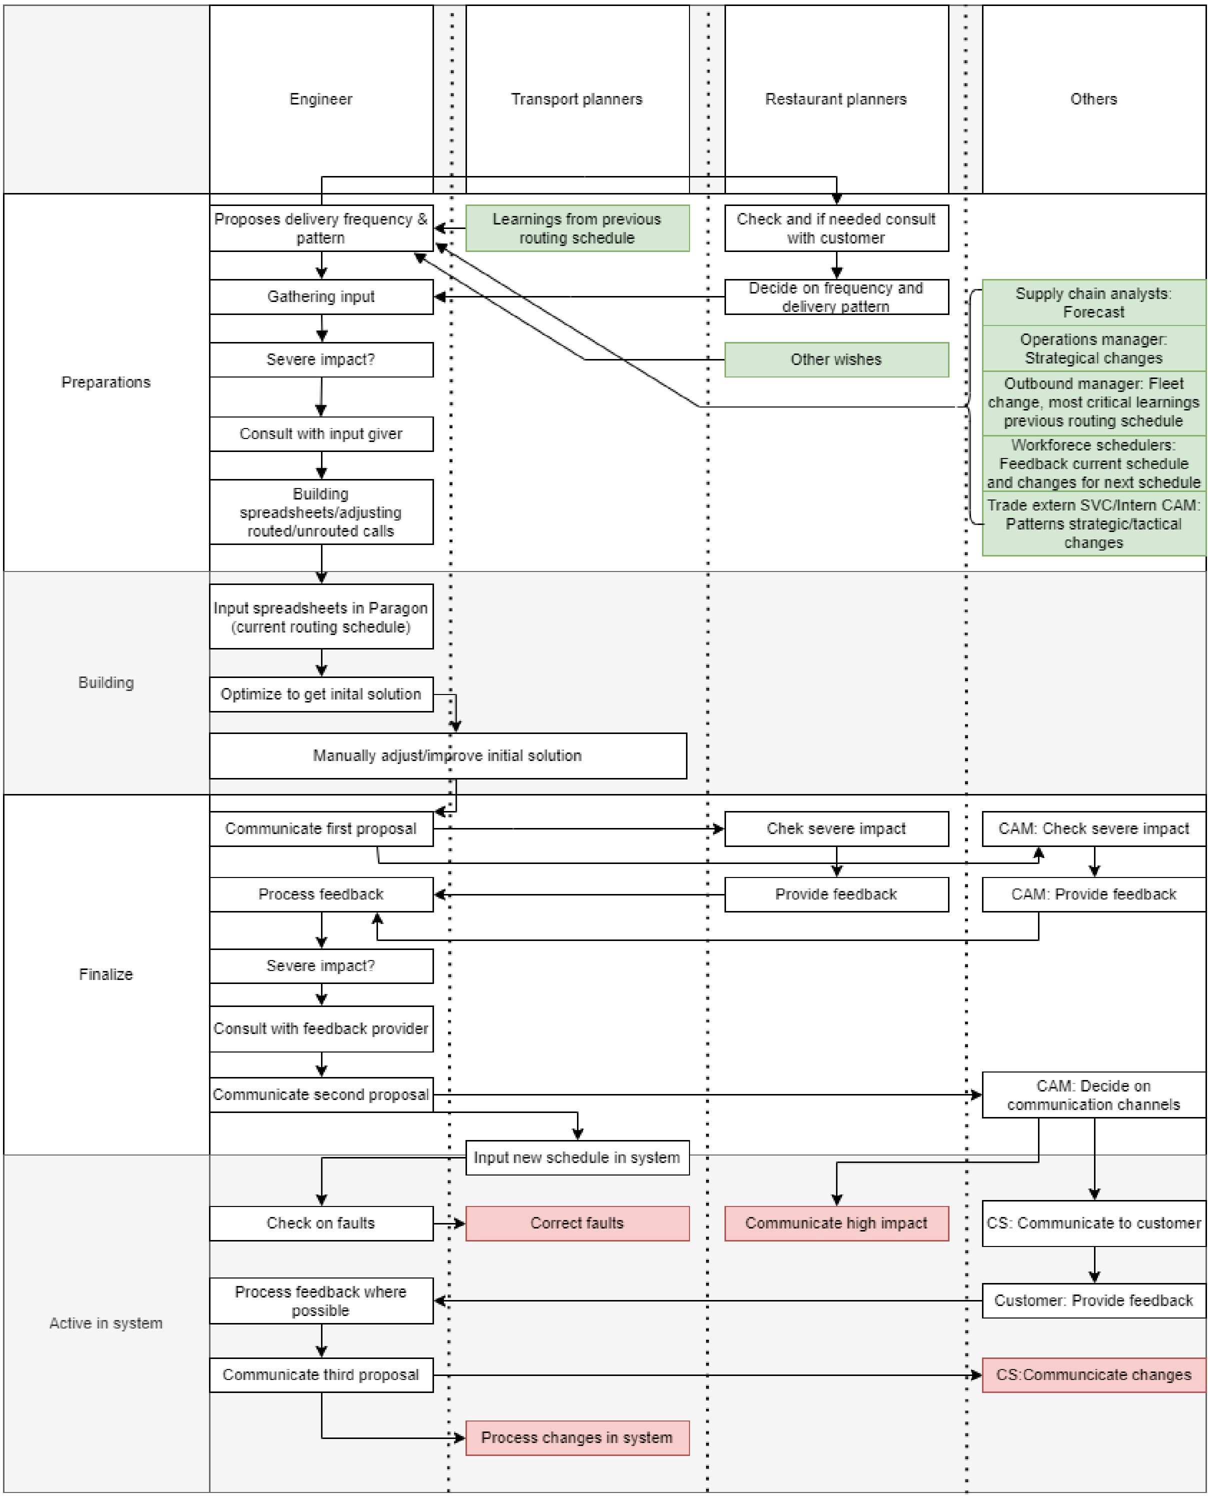

Figure 7: Process of developing a new tactical routing schedule

Simulation studies

The transportation processes as described in the two paragraphs above, are executed by the operational transport planners. The tactical routing schedule is being made by the engineer together with the transport planner. The last years, HAVI uses a transport software system to deal with strategic questions. An example of such a question is: ’How would a new business partner fit in the transport network?’ or ‘Where to locate a new distribution center (DC) if the current DC would be too small in the future?'. This software system, called Paragon, optimizes the routing schedule given deliveries with a fixed volume. In the last year, Paragon is also used to plan the holidays and the December month in which McDonald's has a special discount for every day, which results in enormous volume fluctuations. Since the beginning of 2020, Paragon is used to propose a new tactical routing schedule. The quantities and delivery days are determined in spreadsheets, based on the forecasted sales and the capacity of the storage locations of the restaurants. The restaurant planners had a lot of feedback on the proposed delivery patterns. In practice the opinion of the restaurant planners overrules the analysis made, because HAVI knows that the analysis is based on insufficient data. Paragon already supports in creating routing schedules but is not able to determines the delivery patterns.

Transportation software

Paragon is the optimization software, HAVI uses to optimizes their transportation planning. HAVI has used Paragon for strategical simulation studies and is starting to use it for their tactical transportation planning (i.e., proposing tactical routing schedules). Paragon splits the input in four different levels, (i) master data, (ii) routing data, (iii) input for a specific optimization and (iv) the objective settings. The input is too extensive to cover completely, but we will state the most important.

Master data

Within the master data we define all data that is likely to be equal for all optimizations. In the master data all DCs are defined and all possible truck types to be used. We also define the possible loading rates (i.e., a fixed loading time plus a variable time per delivery unit) from which we can choose later on. We define the possible driver types. We also define all customers. For every customer we define the location, opening times, loading rates, possible vehicles to use to delivery this customer. In this research, we only use one type of truck (i.e., a combination truck with a capacity of sixty roll containers). We also define the minimum time that has to be between consecutive deliveries. Within the master data a matrix is constructed that determines the distances between pairs of customers. A map of the Netherlands is used on which areas are drawn, which apply rush hours and speed adjustments. Within the master data for every pair of customers, the expected driving time is calculated. While the routing process is running, the map with adjustment speeds is considered to calculate the precise driving time based on the time of the day.

Routing data

In the routing data we define data that is more likely to vary between multiple optimizations. However, in this research all optimizations are performed with the same settings. In the routing data we partly add new data and partly add information to the data stored in the master data. We add information to the DCs, for example the possible times to load and unload. We define the number of trucks that can be used from each type of truck. As said, we only use one type, and in the routing data we define that we have six trucks of that type. We define the type of products that we use. In our research we use two products. With simple rules we tell the software that some products are not compatible with other products (i.e., they may not be delivered in the same truck). We use this in our experiments to separate regular customers, from some dummy customers. Furthermore, we define the shifts that can be worked by drivers. In this research, we have only defined one shift, although HAVI operates with two shifts a day. However, we have set the possible departure times at the DC such that drivers only depart within the shift times used by HAVI. In the routing data we also tune the algorithm of the optimizer. Many settings can be changed, which often are a trade-off between result and running time. An example is how many customers to include in a k-swap, or how many different options to try before determining the best swap(s).

Input optimization

The input given for a specific optimization is uploaded via csv-files. The most important data are the calls (i.e., the stops that have to be performed in the optimization). For every call, we state many attributes of which the most important ones are: the customer, whether it is a delivery or a collection, the volume, the delivery window, which product, from which DC, in which route, and execute by which driver type and vehicle type. We see that we also include the route number. Therefore, what we input, is actually a routing schedule stating the routes with all necessary data. Some data is being fixed throughout the optimization and some is only used for the upload. For example, the route and DC are only used while uploading to initiate the routing schedule. Other data is being fixed also while optimizing, for example, the customer and the delivery window.

Objective settings

We can choose the objective of the routing optimization. We can choose to minimize either the distance traveled, or the duty time needed, which includes driving time, loading and unloading, waiting time and extra time defined for administration. Because the customers of HAVI often have time windows with gaps in between (e.g., no delivery during lunch time), minimizing distance traveled often results in more waiting time within the schedule. Therefore, we choose to minimize duty time. Furthermore, we can choose to apply clustering, which means that the distance and time from the DC to the first customer and from the last customer back to the DC are excluded from the objective value. This way, the optimizer constructs the routes such that the customers within a route are closer to each other. In the transport requirements in Section 1.2 it is stated that the routes should be constructed such that when volume increases or decreases, the transport planners can adjust the routes without too many changes. The experience is that this task is easier when routes are created such that customers within a route are close to each other. Therefore, the objective in the RVRP solver is to minimize the duty time, excluding the times travelling to the first customer, and travelling from the last customer back to the depot, this objective is called ‘clustering based on duty time’.

2.3 Order Process

For each location, there are fixed delivery days in the week. The order must be placed before 10:30 am on the day before delivery. So, for a delivery on Wednesday morning the order must be placed before Tuesday 10:30 am. Orders for deliveries in the afternoon must be placed before 3 pm the day before.

The sales forecast for the near future (i.e., at least covering the coming few deliveries) is filled in by the restaurant. The system of HAVI, HAVI Core, proposes an order that will be sufficient to meet demand between the delivery moment and the next delivery moment and bring stock back to safety levels (i.e., JIT delivery). Eight restaurant planners, that are responsible for a region of McDonald’s locations, can adjust those orders. They can choose to pull volume from the subsequent delivery to the first delivery. The two main reasons for doing this are (i) overcoming uncertainty (i.e., because of some circumstances, the safety stock level is probably not sufficient), and (ii) removing the peaks in delivery volumes. The former is to deal with exceptions. The focus in this research is on the latter.

2.4 Delivery patterns of McDonald’s

The delivery patterns of McDonald’s customers are largely influenced by an earlier attempt of reducing the delivery frequency. In 2013, the average delivery frequency per store was 5,5 deliveries per week. To minimize costs, the delivery frequency is lowered in the years after. The average delivery frequency was 4,5 deliveries per week in 2015. In 2019 the average delivery frequency at 4,6 deliveries per week. Because of the decreasing delivery frequency and meanwhile an increasing volume per store, the volume per delivery is increased. Because consumption is highest in the weekend and restaurants also prefer a delivery right after the weekend, most deliveries that were removed were Tuesday deliveries. Most restaurants that have a delivery on Monday but not on Tuesday want a delivery on Wednesday. This created peaks in delivery volumes, which are still present. Nowadays, HAVI copes with peaks on Monday, Wednesday, Friday and Saturday, and quieter days on Tuesday, Thursday and Sunday.

The delivery patterns are fixed, because then the stores can consider the delivery in their staff planning. Because of this staff planning, amongst other reasons, it is often the preference of the store to have all deliveries in the morning or all deliveries in the evening (they often have one or two employees, who are responsible for the delivery). Next to these preferences, we cope with legislations of loading and unloading in the cities, which result for example in stores having all their deliveries in the early morning. These preferences and the overall way of thinking of HAVI and their customers, requires a standard routing schedule that is repetitive every week. The impact of applying all available data to make a routing schedule that is not repetitive each week would be the next step in research. For this research this is out of scope. It is important to have a good standard routing schedule because adjustments in the operational planning are minimized and workload can be flattened within the week.

2.5 Conclusion

In this chapter we answer research question 1: ‘*How does the current situation look like?’*. The first sub question is: *How does HAVI execute the transport optimization?* First, HAVI has a standard weekly routing schedule that is repetitive over the weeks. This is the tactical planning, which is used to change the routing schedule when the circumstances change, including seasonality changes. The standard routing schedule is the basis for the operational planning. When exact volumes are known, the routes are adjusted for the following day. For strategic questions, simulation studies are executed. Paragon, the optimization software, is used to create hypothetical routing schedules. Last year, Paragon is also used more intensively to help creating the standard routing schedule. It is experienced that the proposed patterns, which are determined in spreadsheets, are not accepted by the restaurant planners. The potential of Paragon is not yet utilized.

The second sub question is*: How does the order process of HAVI look like?* Given the delivery moment, there is a fixed order moment. HAVI’s system, proposes the order. The restaurant planners can adjust the orders to remove the peaks in delivery volumes.

The delivery patterns result in peaks, which are the result of earlier projects with the goal of minimizing the number of deliveries per restaurant. Improving the tactical planning, flattens the workload peaks and minimizes the number of adjustments needed in the operational planning.

3. Literature study

In this chapter, we answer research question 2: ‘*How is volume or workload balancing dealt with in the models in literature (i.e., PRPs and IRPs)*?’. In Section 3.1 the processes of HAVI are related to known models in literature. These models are explained in Section 3.2 to 3.5. After elaborating on the models, they are again related to the processes of HAVI in Section 3.6. In Section 3.7 we discuss various solution approaches used in literature and discuss which aspects are essential for us to consider.

3.1 Literature models related to HAVI processes

HAVI uses a tactical routing plan that consists of routes for seven consecutive days. This routing plan serves as a basis for every single week within a period of months. Every day, the transport planners are planning the exact routes for the following day. At this moment, the quantities to be delivered are fixed. The transport planners’ job is to solve a vehicle routing problem (VRP) with as little adjustments as possible from the tactical routing plan, whilst eliminating inefficiencies. A few times per year, a new tactical routing plan is made to adjust for seasonality and changes in the market. While creating a tactical routing plan, the delivery patterns are not fixed and can be redetermined. The new created tactical routing plan serves again as a basis for seven consecutive days that is repeated every week. In this aspect, this is a periodic routing problem (PRP). While making the tactical routing plan, the quantities to be delivered are not known yet. In the current way of working, a forecast for a relatively high-volume week, is taken as if demand is fixed. Therefore, this literature study also focuses on models that are based on deterministic volumes.

The aim of this study is to research how HAVI can integrate the tactical transport planning and the inventory management of their customers’ locations. This points us into an inventory routing problem (IRP). The IRP is meant to deal with the situation where routing problems and inventory problems are combined into one model.

In the next part of this literature study, we elaborate on the PRP and IRP models and how they are developed. Both PRP and IRP are established after further development of the VRP. Therefore, we first discuss the VRP.

3.2 Vehicle Routing Problem with side constraints

Laporte, Toth and Vigo (2013) describe the vehicle routing problem as a problem consisting of designing least costs delivery routes through a set of geographically scattered customers, subject to several side constraints. The vehicle routing problem is one of the most studied combinatorial optimization problems. Golden, Raghavan and Wasil (2008) say vehicle routing may be the single biggest success story in operations research. The type of constraints that are added to the model, results in many special cases and varieties of the vehicle routing problem. In our research we must deal with the following:

- Capacity constraints of the vehicles.
- Capacity constraints of the inventories at the customers’ locations.
- Constraints regarding the perishability of the goods delivered.
- Uncertainty in demand at the customers’ locations.
- Delivery time windows.
- Periodic schedule of one week.
- Largest customer VMI managed. The rest of the customers place their orders themselves.

3.3 Capacitated vehicle routing problem

The capacitated vehicle routing problem (CVRP) is the most common extension of the standard VRP. In the CVRP, a fleet of trucks, possibly with different capacities, is used to carry out the routes. The CVRP was for the first time formulated in a mathematical way in 1959 by Dantzig and Ramser. They called it “The truck dispatching problem”. In this research, we denote it by CVRP or even VRP for consistency. Dantzig and Ramser take the traveling salesman problem (TSP) as starting point. A TSP is a special case of the VRP in which only a single truck with enough capacity is used to deliver all customers in one trip and the shortest route had to be found visiting every location exactly once.

The VRP as described by Dantzig and Ramser (1959) deals with a single bulk product (i.e., gasoline). Given is the demand per station and the shortest route between any two points within the system. We elaborate on their method with an example in Appendix IV. Dantzig and Ramser (1959) are the first to formulate the VRP in a mathematical way. Their method does not necessarily lead to the optimal solution, but the methods used are, although not totally straightforward, relatively easy to follow. In Laporte (2009), fifty years of development of the VRP is discussed. Most algorithms used in practice are heuristics. Lenstra and Rinnooy Kan (1981) state that VRPs are NP-hard problems. Therefore, exact algorithms can only solve small instances. Laporte (2009) says, instances of up to one hundred vertices can be solved in reasonable time.

3.4 Periodic Routing Problem

In 1974 the PRP is for the first time introduced by Beltrami and Bodin (1974). A PRP is a special case of a VRP. The PRP is defined by Beltrami and Bodin (1974) such that a set of customers requiring a fixed number of deliveries in the planning period of days, finding the best allocation of customers to the schedules, minimizing the costs of visiting the customers. Since a PRP, or also called PVRP, is a special case of a VRP, mostly heuristics and rarely exact algorithms are used. Lahyani, Khemakhem and Semet (2015) say that in a PVRP all the input data is available at the beginning of the planning period. In a PRP, the decision maker must decide when to visit each customer and how to construct the routes to deliver the customers. For each customer, a set of delivery patterns exists from which a pattern is chosen, as also described by Carotenuto, Giordani, Massari and Vagaggini (2015). Choosing the patterns first and without considering routing, this results in a VRP per day that must be solved. The decision maker chooses the delivery patterns and routes separately, which results in a suboptimal solution. This could be related to the fact that delivery quantities are being fixed after determining the delivery frequency f_i for customer i. Francis, Milowitz and Tzur (2008) state that in the PVRP literature it is assumed that a fraction $1/f$ of the total demand must be delivered to customer i each visit. Campbell and Wilson (2015) are more generic, stating that for every customer i there is a set of feasible visit options Λ_i (i.e., delivery patterns). They state that if a product is to be delivered or collected, the quantity will be known and will be fully satisfied by a single vehicle. Therefore, the delivery quantities are a consequence of the delivery pattern.

HAVI schedules the deliveries different from how deliveries are scheduled within a PRP. Within a PRP, the quantity to be delivered is dependent on the delivery pattern that is chosen. A pattern only states the delivery days and not the delivery times. However, for HAVI, delivery quantities are dependent on delivery times. The volume consumed within the period of a day is significant relative to the average delivery size. Some locations have a small storage, which does not allow us to define delivery quantities independent of the delivery time. Because we have to include the storages into the model, we look at IRPs in the next section.

3.5 Inventory Routing Problem

The basis of IRPs is clearly discussed in Cordeau et al. (2006). They define an IRP as an extension of the VRP in which inventory control is restricted to ensuring that no stockouts occur at the customers. IRPs arise in VMI environments. The vendor can choose the timing and size of deliveries to a set of customers. The vendor agrees to ensure that its customers do not run out of product. The vendor also decides which routes to travel to deliver the customers. Inventory costs may or may not be included in the optimization. IRPs are also NP-hard (Adulyasak, Cordeau and Jans, 2013) and at least as difficult to solve as PRPs. Therefore, they are mostly approached with heuristics and only exact solutions are used in small instances.

Bell et al. (1983) are the first to describe an IRP problem. They explain the benefit of an IRP by use of an example that is often referred to by others in literature. We elaborate on this example in Appendix IV.

One fundamental difference between PRP and IRP is that in PRP the size of deliveries is known in advance, whereas the delivery size in an IRP is a decision variable. It is essential for an IRP to include long-term optimization because optimization over the planning horizon only, results in postponing as much demand as possible to the period after the planning horizon (Cordeau et al. , 2006).

3.6 Relating to HAVI

At the operational level, the transport planners are solving VRPs. At the tactical level HAVI has its own way of working in which we see some similarities with PRPs and IRPs as well. The aim of IRPs is to integrate the transport planning with inventory management. Both PRP and IRP take the delivery day per customer as a decision variable.

For the operational planners this is not the case because the set of customers to be delivered is fixed on forehand. For making the tactical routing plan, the set of customers to be delivered at each stage is free to be chosen. Although the decision space for HAVI is the same as for an IRP, there are some fundamental differences. After the tactical routing plan is made, this routing plan must serve as basis for multiple consecutive weeks that can differ significantly from each other. The timing of the deliveries in the basis must be sufficient for every single week in the period when the tactical routing plan is active. The decision, when to deliver each customer is dependent on the volume fluctuations over the weeks. Furthermore, it is necessary to do the inventory stock calculation right before and after delivery, to avoid stockouts or exceeding capacity. In some models the sequence of operations is fixed. For example, in Bertazzi and Speranza (2013), it is assumed that delivery always takes place before consumption. They use the logic from the lot sizing literature. These models are not applicable for us, since delivery can take place after some part of consumption of that day has taken place.

One of the most important aspects is to adjust the delivery quantities based on the delivery times. We did not succeed in identifying applicable literature that addressed this. The dependencies are illustrated in Figure 8.

The complexity arises because the quantity delivered to a customer is dependent on the delivery time, which depends on the routes constructed. The routes depend on the set of customers that are chosen to deliver on a specific day. However, the decision to deliver a specific customer depends on the rest of the customers, with whom it can form a route. Even a set of customers that can form a route is complex to define on forehand because the set of customers that fit in a route, depends on the quantities that must be delivered to each customer, which is again dependent on the delivery time. Because of this circular reasoning, it is hard to find the order of steps that must be performed to find a good solution.

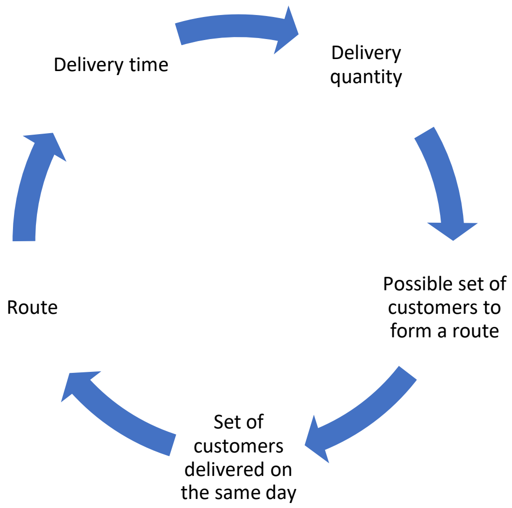

Figure 8: Complexity of delivery quantities based on delivery times

Summarizing, there are four aspects that make this research unique:

1. Creating the tactical routing schedule, the transport and inventory is integrated in a single optimization phase, instead of a decomposition by the two disciplines.
2. The tactical routing schedule is combined with an operational policy using IRP in which the operational IRP is bounded by the tactical routing schedule.
3. The tactical PRP considers that volume to deliver is dependent on the delivery time.
4. The tactical PRP considers that volume balancing can be adjusted using the operational policy.

3.7 Solution approaches

We now discussed two main classes of models that are used to deal with multi-day routing optimizations, namely PRPs and IRPs. In this paragraph we elaborate more on the various approaches used in the literature, and especially zooming in on the aspects that can help us with balancing the workload.

PRP

Bertazzi and Speranza (2013) explain that within a PRP, the quantities of delivery are being fixed now that the days of delivery are determined. Baptista, Oliveiro and Zúquete (2002) elaborate on various solution approaches for the PRP used in earlier research. They conclude that just a few solution procedures are described in literature. One approach first makes routes, following by assigning those routes to the days of the planning period. The second approach first selects the delivery days for the customers and then creates the routes, which is equal to solving VRPs, because in a PRP delivered quantities are fixed when the delivery pattern is chosen.

Russell and Igo (1979) assign customers sequentially to delivery days. Based on customers already assigned, more customers are being assigned to a day. The chosen delivery day is based on different spatial factors, such as closest customer, which is already assigned to a day, average distance and variance of distance to customers in a cluster (i.e., which could form a route) on each day.

The characteristic of a PRP, which states that delivery quantities are known when delivery days are chosen (Campbell and Wilson, 2014), does not satisfy our needs. Because we believe in an approach in which we can shift volume in the same phase as we are optimizing the routes, we focus in the rest of this chapter on the IRP model.

IRP

Before we elaborate on the various approaches to solve an IRP, we will narrow our scope down by using a classification of the IRP as stated by Bertazzi and Speranza (2012) and describe some of the characteristics and often made assumptions of various IRPs. Bertazzi and Speranza (2012) classify the IRP based on the decision space as follows:

1. *Decisions over time only*. The timing and quantities of the deliveries must be decided, while the routes are given.
2. *Decisions over time and space*. The timing and quantities of the deliveries must be decided as well as the routes to be traveled.

Bertazzi and Speranza (2012) focus on IRPs with decisions over time only. In our research we also focus on an IRP with decisions over time and space. Both the timing of the deliveries as well as the routes are crucial in our research to create a more efficient solution.

Bertazzi, Savelsbergh and Speranza (2008) say that the basic characteristics that have a large influence on the formulation of the IRP are the following:

- the planning horizon can be finite or infinite.
- inventory holding costs may or may not be considered.
- inventory holding costs may be charged at the supplier only, at the supplier and the customers, or at the customers only.
- the production and consumption rates can be deterministic or stochastic.
- production and consumption take place at discrete time instants or take place continuously.
- production and consumption rates are constant over time or vary over time.

the optimal delivery policy can be chosen from among all possible policies or must be chosen from among a specific class of policies

In our research, we focus on an IRP with a finite horizon, in which inventory costs are not considered. Production is not within the model and consumption rates are assumed to be deterministic. Consumption takes place at discrete time instants and vary over time. The optimal delivery policy should be such that the final inventory levels equals the inventory levels at the beginning of the planning horizon.

Since the IRP is mentioned in literature, technology has evolved with huge steps. Data is being transferred between vendor and retailer for on a larger scale. The vendor is often able to access the data needed to have a clear view (after some statistics possibly) on the pattern of the inventory during a day of each retailer. However, most models in literature are still based on the assumption that consumption takes place at one specific moment of the day. Bertazzi and Speranza (2013) assume that in an IRP several operations takes place in a specific order. The order of the operations each day is assumed to be 'delivery-consumption-calculation of inventory'. The choice of the order in which these operations take place is always of great influence of the mathematical formulations. Bertazzi et al. (2008) elaborate on the example constructed by Bell et al. (1983) as illustrated in Appendix IV. Bertazzi et al. (2008) assume that delivery takes place before all consumption of that day. They say that the discrete time model with the assumption of delivery taking place before consumption of that day, is too restrictive for many environments. For example, environments where products are used continuously, as it is in our case study. There are several cases in which consumption is not assumed to be consumed at one specific moment of the day. In these cases, either the assumption is made that demand is constant (e.g., Coelho, Cordeau and Laporte, 2012; Archetti, Bertazzi, Hertz and Speranza, 2012), or that demand is stochastic (e.g., Trudeau and Dror, 1992). With a constant demand, balancing workload is not necessary. For us, a stochastic demand is out of scope, because the software that is available only takes deterministic input.

A few more assumptions are often made, which are worth mentioning before we elaborate on the solution approaches. Initial inventory levels are often assumed to be known. They are often either zero or equal to maximum capacity. In Bertazzi et al. (2008), both cases are used as well as the case in which the initial inventory levels are decision variables. Another common assumption in an IRP is unlimited shelf life. Soysal et al. (2015) say this is one of the main obstacles for the application of the basic IRP models the food sector. They do include the shelf life themselves in a model with stochastic demands. They penalize the expected waste per customer per period. In their research they build an algorithm, which has to solve the routing problem. For our research the algorithm of the solver used is out of scope, and we deal with deterministic volumes when creating a tactical routing schedule. Therefore, we do not have expected waste in the tactical routing schedule. Waste has to be considered in our research in the operational planning only.

Within IRP models, different approaches are used to determine the delivery quantities. There are several approaches that determine the delivery quantities independent from the routing process. The simplest forms are the ones in which the inventory of the customer is always fully utilized after delivery. Fisher, Greenfield, Jaikumar and Kedia (1982) applied this logic in the gas oil industry where the quantities are fixed such that the tank for gas oil is fully filled after delivery. Bertazzi, Paletta and Speranza (2002) use the order-up-to-level policy to determine the quantities. Every time a retailer is visited, the quantity of each product delivered by the supplier is such that the maximum level of the inventory is reached at the retailer (Bertazzi et al. 2002).

In Dror, Ball and Golden (1985) and Dror and Ball (1987), the delivery quantities follow from the decision of delivery days instead of the delivery quantities being a decision variable. In this aspect this is the same as in a PRP, except that the delivery patterns are not chosen from a set of patterns that are created in advance. Some of the earlier researches, such as Anily and Federgruen (1990), use EOQ formulations to determine the delivery frequency. They construct regions of customers that are geographically close to each other. Each region is a subproblem. A route is made to visit all the customers in the region. With an EOQ based formula, it is determined what quantity is economically most efficient to deliver every time that route is driven. The drawback is that every customer within a region must get the same number of deliveries.

All the approaches as described above determine the delivery quantities, either after the routes are being fixed or the routes are not even considered yet. These approaches have laid the foundations for further research, but these approaches will not be used in our research because we are aiming for integrating the determination of delivery quantities and constructing the routes into one phase.

Moin and Salhi (2007) state the categorization of the solution techniques as follows: "Solution techniques for the IRP can be classified into two categories namely the theoretical approach, where the derivation of the lower bounds to the problem is sought and a more practical approach, where heuristics are employed to obtain the near-optimal solutions.". In the theoretical approach, most papers split the problem into two phases: the inventory and the traveling salesman problem. They choose for a two-phase approach in which they start with one of the problems, followed by the other. Most algorithms iterate between finding the routes and solving the corresponding inventory problems until the stopping criterium is met.

Since we are doing a case study, we want to obtain a near-optimal solution instead of a lower bound. Therefore, we focus on the practical approach where we probably need a heuristic to find a near-optimal solution. However, in the practical approaches, not only heuristics are used. Often mathematical formulations such as integer problems (i.e., ILP or MIP for example) are used for sub-problems. Campbell and Savelsbergh (2004b) use an ILP model to determine the delivery days and quantities. Lee, Bozer and White III (2003) use an MIP to determine the optimal inventory levels per route. Bell et al. (1983) use an MIP to determine, after the routes are created, which routes are driven each day.

Also, within the practical approach, many researches choose for a decomposition of the IRP into sub problems (i.e., multiple phase approaches). Hulshof (2008) states that a decomposed approach (i.e., routing problem and inventory problem separated) is better than an integrated approach. The reason therefor is that the integrated approach increases the problems complexity and therefore requires prohibitively large computation times. Hulshof has researched a case study at ORTEC in which a two-phase approach is used in which first the orders are created based on forecast, and second the VRP is solved. In Dror et al. (1985) and Dror and Ball (1987), first the delivery days are determined for each customer. The delivery quantities follow from the fact that the inventory of the customer must be full after delivery. This results in a VRP. Bell et al. (1983) first create a set of routes. Afterwards they choose which routes are driven each day. Lei, Liu, Ruszczynski and Park (2006) use a two-phase approach in which the first phase is about determining the timing and quantities of the deliveries. The second phase solves an associated consolidation problem to make more efficient routes than the direct shipments that were created in the first phase. Campbell and Savelsbergh (2004b) use a two-phase model. In the first phase they determine the delivery days and quantities. In the second phase the distribution plan is created. Lee et al. (2003) also use a decomposition into a VRP and inventory control. The vehicle routes are generated using a simulated annealing method. Given a set of routes, they use linear programming to determine the optimal inventory levels per route. Based on these inventory levels, they reschedule the routes after which they solve the linear program again. This process is iterative with a predefined number of iterations.

Earlier this section, we stated the basic characteristics as mentioned by Bertazzi et al. (2008). The first characteristic is about the planning horizon. For infinite horizon models, often a selection of customers is set as a region. The regions are chosen such that the demand of the customers is roughly equal to a full truck load. Infinite horizon models only cause asymptotic results, which can be misinterpreted in practice according to Moin and Salhi (2007). In our research, we focus on finite horizon models. When dealing with a finite horizon, however, we cannot ignore the long term. Ignoring the long term makes it possible to postpone as much as possible to the period after the planning horizon, which has a negative effect on the costs in the future. Cordeau et al. (2006) say that the long-term effect of short-term decisions needs to capture the costs and benefits of delivering to a customer earlier than necessary. Fisher et al. (1982) and Bell et al. (1983) pioneered in an approach of incorporating long-term costs in a short period planning period. They assign an expected cost per unit for a delivery beyond the planning horizon. Dror et al. (1985) and Dror and Ball (1987) are the first to study the effect of short-term decisions on the long term. They first calculate the optimal delivery day. If this delivery day is within the short term, the customer is visited in the short-term planning period. If the optimal delivery day falls outside the short-term planning period, they apply a future benefit when they do deliver the customer within the short-term planning period. Bard, Huang, Jaillet and Dror (1998) introduced the approach in which two different horizons are used. The planning problem is defined for the first two weeks, but only implemented for the first week. After the first week, the problem is defined again for the coming two weeks, of which the first week is implemented. Campbell and Savelsbergh (2004a) also use the principal of the rolling horizon. The distribution plan is created for one month, but only implemented for the coming few days. Hulshof (2008) also distinguishes between two horizons. Hulshof (2008) mitigates the peaks in a planning horizon of a whole year. However, the routes are only created for customers that need a delivery in the short term using a planning horizon of a few days.

Research question 1b is about balancing the workload. Waller et al. (1999) say that many suppliers are attracted to VMI, because of the mitigation of the uncertainty of demand. However, as in our case study appears, even in a VMI environment, the uncertainty of demand is not the only factor that influences the peaks in workload. The sales patterns of the customers contribute the most to these peaks. We can say that VMI mitigates demand peaks because the data is more transparent (Andel, 1996). The vendor can forecast the volume of the coming period based on actual sales of the previous days. In a customer managed inventory environment (CMI), the volume of the coming period is more uncertain and can be based on actual orders.

We now discuss how we can deal with balancing the workload. Campbell and Savelsbergh (2004c) have discussed efficient implementations of insertion heuristics to handle situations where the delivery quantity must lie between a lower and an upper bound. In this implementation phase they use estimated delivery costs for a large set of possible clusters. They define the customers that must be visited in the short-term planning horizon as must-go visitors. In the basis, a cluster is formed of a group of must-go customers. Additional customers are added based on geography and capacity. This will not directly balance the workload over all customers aggregated, but it can be used to balance the workload. Hulshof (2008) uses the separation of the planning period in a short-term and long-term planning horizon, as described in the previous paragraph, to deal with volume balancing. Hulshof (2008) developed an algorithm that selects delivery days for every customer. The algorithm selects the customers that require a delivery in the short-term planning period. Following by adding customers, that do not necessarily require a delivery in the short-term planning period, but who can receive a relatively large delivery compared to their capacity. Simultaneously they also tradeoff the extra costs of adding the extra customers to the distribution costs. The main concern with volume balancing is balancing the volume over a year. Additionally, volume balancing on short term is incorporated. The balancing of the volume on the short term is indirectly done based on balancing workload. To balance the workload, the number of clusters visited on a day has a lower and an upper bound. Clusters are made using seeds. The seeds are located such that the sum over all customers of the distances between the customer and the closest seed is minimized. The short-term workload is balanced using a maximum and minimum workload per day. The maximum and minimum workload is dependent on the available capacity of a given day, relative to the total capacity of the complete period. This ratio determines the optimal volume delivered on that day. A factor β is used to be added and subtracted from this optimum to make the problem feasible. The problem is solved using an ILP problem in which clusters of customers are selected per delivery day.

In the previous paragraph we showed how workload balancing can be incorporated within a heuristic. Volume balancing can also be accomplished by including inventory costs. Bertazzi et al. (2008) start with a simplistic model, after which more complexity is added to study the effects of some characteristics. For example, they add inventory costs, to study how the solution changes. When studying their examples, we learn that when inventory costs at the customers' locations become higher, the goods are stored at the customers' locations no longer than needed. With higher inventory costs, the system looks like a JIT system in which every customer is being delivered every day. The examples are too simplistic to be translated to our case study because demand and production are assumed to be constant. We can use the same logic and intuitively argue that when high inventory costs at the supplier are used, the goods will not be stored longer than necessarily at the supplier's location. When a constant production rate is assumed, this results in an even distributed volume every day, which balances the workload. Lei et al. (2006) include production into their IRP, which becomes a production, inventory, and distribution routing problem (PIDRP). In the first phase of their approach they solve a restricted problem in which they keep all the constraints except that the transporter routings are limited to direct shipments. In the second phase they improve the routing. The approach as used by Lei et al. (2006) could help us with workload balancing. However, we do not use their approach because we believe in an approach in which the timing and sizing of the deliveries is decided upon in the same phase as making the routes.

We close this chapter by addressing an aspect that is almost neglected in literature, namely, time dependent delivery quantities. Bertazzi et al. (2008) and Campbell and Savelsbergh (2004b) study the time dependent delivery quantities. Given a route (i.e., sequence of customer visits), they determine the optimal delivery times, to maximize the total quantity delivered to the customers. The maximum quantities that can be delivered are dependent on the storage left at the customer, and the discharging time. For our research, the discharging time is out of scope when determining the possible delivery quantity. Bertazzi et al. (2008) use these time dependent delivery quantities to maximize the utilization, given the route. We do not aim at utilizing the capacity of the truck when routes are given, but we want to use the time dependent delivery quantities to make the routes simultaneously with determining the delivery quantities and determining the set of customers that will form the route.

3.8 Conclusion

This chapter answers research question 2: '*How is volume or workload balancing dealt with in the models in literature (i.e., PRPs and IRPs)?'*. The first sub question is: *Which solution approaches are used to solve the models?* The tactical and operational transport planning at HAVI can be described by an IRP model. The VMI environment HAVI works in, has benefits regarding accessibility to data, which is needed to feed accurate data into the IRP. We focus on practical approached IRPs with decisions over time and space. The sequence in which different operations (for example, delivery-consumption-calculation of inventory) takes place is highly influencing the mathematical formulations of the model. The assumption of such a sequence is, especially in the food industry, a reason why the IRP is hardly applied, as well as the assumption of unlimited shelf life.

In earlier research, delivery quantities and the routes created are often determined independently. We were not able to identify the literature that jointly optimizes the routing problem and the inventory control. The approach is to decompose the problem into subproblems: inventory control and routing. These are considered in several phases one after the other, sometimes with multiple iterations. The reason of this research is to integrate both the disciplines. In the current way of working this decomposition leads to inefficiencies.

The second sub question is: *How do these approaches deal with volume or workload balancing?* Workload balancing can be included, by first selecting customers that require a delivery in the short-term planning period. Thereafter customers that not necessarily require a delivery in the short-term planning period, can be added to the routes based on the workload of the already planned customers/routes whilst including geographical criteria. Also, including a constant production rate in combination with a high inventory cost at the suppliers' warehouse, forces the volume to be delivered in a constant rate resulting in mitigation of peaks.

Time dependent delivery is addressed sparsely in literature, and where discussed, it is used to utilize the truck capacity given a route, instead of optimizing the set of customers that form a route. In Figure 8 the complexity of the timing of the deliveries is visualized. Because of this complexity this research will use a method in which (i) the routes or (ii) the delivery quantities are not based on the other aspect being determined in a previous phase. We build a model in which they are integrated in a single phase.

Concluding, for the creation of the tactical routing schedule, we must use the insight and the same kind of restrictions as with an IRP. The mathematical formulation of the model should be adjusted such that shelf life restrictions are included. This makes the model suited for the food industry. To get rid of the dependencies of decisions in different disciplines (i.e., transport and inventory) we must formulate the model such that the transport and inventory are dealt with simultaneously.

4. Model building

In this chapter we answer research questions 3, 4, 5, 6b and 7. These research questions are as follows:

Research question 3: '*How can the requirements be translated into restrictions/input that can be used in routing optimization software?'*

Research question 4: *'How can workload balancing be applied in the tactical routing process of HAVI?'*

Research question 5: *'How do we build the model to optimize the transport and inventory integrated?'*

Research question 6b: '*How to compare the results of the experiments?*'

Research question 7: *'What operational policy should be followed after implementing the new tactical routing schedule?'*

In Chapter 2, we described the current situation, for which we will develop a model such that the routing and inventory management are integrated into one system. In Chapter 3, we elaborated on the literature, which provides us with models that are used in optimization studies for routing and inventory management.

We concluded in Chapter 3 that the assumption of unlimited shelf life is one of the main obstacles in the basic IRP model. The assumption is not valid in the food sector and not at HAVI. The different studies in literature mostly apply a multiple phase approach. In this chapter we build a model in which we use a single phase in which the delivery pattern and delivery quantities are being optimized simultaneously in the tactical routing schedule.

The main characteristics that we must include into our tactical model are:

- Delivery quantity is dependent on the exact timing of the delivery, not only the delivery day.
- Multiple groups of products are considered in storage as well as shelf life.
- The delivery patterns are optimized simultaneous with optimizing the routes.
- Workload balancing should be considered.
- We should steer on minimizing the workload on more expensive days.

We first build a model to improve the tactical routing schedule. Afterwards, we build a model for the operational policy. For the tactical model, we discuss in Section 4.1 the process of optimizing the routing schedule. We follow by discussing why and how we are going to optimize the delivery patterns and quantities simultaneously in Section 4.2. In Section 4.3, we discuss the input that we use. We continue with deducing lower bounds for the storage capacities in Section 4.4. Followed by how we determine the delivery windows for every delivery unit in Section 4.5. In Section 4.6 we use the assumption that the current routing schedule is a feasible solution, to find higher lower bounds. Finally, we build the model consisting of an allocation procedure and two improvement heuristics in Section 4.7. We use Section 4.8 to verify and validate the model we have built. In Section 4.9 we build the operational policy that is used to execute a sensitivity analysis.

For this research to be beneficial for HAVI, they first need to know the capacities of the customers' storage locations. If capacities are already known, Sections 4.3, 4.4 and 4.6 can be skipped. In this research, we use an RVRP solver, which performs better when an initial solution is provided. In the case that there is no initial solution available or an RVRP solver is used who does not need an initial solution, Sections 4.6, 4.7.2 and 4.7.3 can be skipped. In those cases, the RVRP solver is used to construct routes out of the delivery units belonging to customers. However, the rest of this research is formulated for the scenario in which we use an initial solution, because that is the case for HAVI. The experiments and results are also based on the use of an initial solution.

4.1 Process of routing schedule optimization

Because we want to jointly determine the delivery patterns and the delivery quantities in our model, which we did not find in models in literature, we need a revolutionary approach. To separate the determination of the delivery patterns and delivery quantities, we step back from the idea that we create deliveries and that those deliveries should be assigned a quantity. Instead we create all delivery units and for each delivery unit we determine when to deliver. From the result, we can see what delivery patterns are created and how many delivery units are assigned to the delivery moments, determining the delivery quantity. Therefore, the output are the delivery patterns and delivery quantities, but these are not the decision variables.

In this section, we discuss the process of optimizing the tactical routing schedule. The model we build is an improvement heuristic. For this research we use an RVRP solver. HAVI uses Paragon as their transportation software, which is able to solve RVRP problems. Other RVRP solvers can be used with all kind of extra restrictions that are necessary for the specific application. Because of the complexity and asymmetric aspects of the routing problem (see Appendix V), optimization often results in a local optimum. By changing the input for the RVRP solver, we help the RVRP solver escape from local optima.

 shows the process of our improvement heuristics and the use of an RVRP solver. In this research, we build the improvement heuristics that changes the input data for the RVRP solver. We then use an existing RVRP solver to optimize the routing schedule. In this research, the algorithm used to optimize the RVRP is out of scope.

Because we build an improvement heuristic and no construction heuristic, we need an initial solution. In this research, we take the current tactical routing schedule as the initial solution for our model. As the current tactical routing schedule, we take the routing schedule HAVI executes the period from the 6th of July until the 6th of September 2020. In this research, we use the term 'current (tactical) routing schedule' which can be replaced by 'initial (tactical) routing schedule'. The volumes we use for the new tactical routing schedule, are also the same as HAVI used for the current routing schedule. Therefore, this research is used to find how the new tactical routing schedule performs compared to the current situation. In the future, for creating new tactical routing schedules, the initial solution should be set equal to the routing schedule that is active at that time, but volumes are forecasted for the period for the next tactical routing schedule.

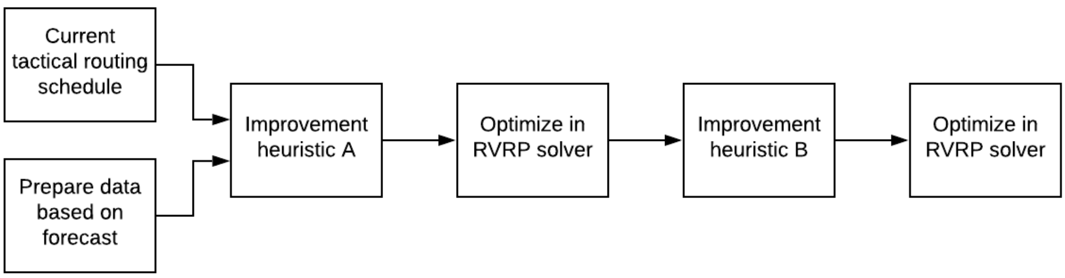

Figure 9: Diagram improvement heuristic

The initial solution is a set of routes for seven consecutive days that determines the sequences of customers within each route, but without any volumes determined. In improvement heuristic A, we use an allocation procedure to allocate all delivery units to the routes. The result is the same set of routes as the initial solution but filled with delivery units. This is the input for the RVRP solver. The RVRP solver is used to optimize the routing schedule. The result of the RVRP solver is a new set of routes in which the same delivery units are routed. However, for improvement heuristic B we only need the set of routes, without the delivery units. Because in heuristic B we allocate all the delivery units again using an allocation procedure. We now have filled the new routes. This is the input for the RVRP solver, which again solves the routing schedule.

Figure 10 shows the process diagram in a more elaborated version. First, we briefly explain the diagram. The rest of this chapter discusses this process regarding the optimization of the tactical routing schedule in more detail. The main input for integrating transport and inventory management is the capacity of the customers' storage locations. The tactical routing schedule will be made with fixed volumes, which are equal to the forecast. Based on this forecast, we make a list of all delivery units that we must plan in the tactical routing schedule. This list is used later multiple times, marked with a star in Figure 10. For a restaurant with forecasted volume of D per week, the list contains D lines for that restaurant. Every line states which restaurant, temperature zone and unique ID corresponds to the delivery unit. For each delivery unit, we determine the delivery window. These delivery windows are added to the list of delivery units so that it forms the input for the RVRP solver. Before sending the input to the RVRP solver we execute heuristic A in which we allocate all delivery units from the list to the routes from the initial solution. The settings in the allocation procedure determine the priorities used, resulting in some shifts to be more utilized than others. We then use the RVRP solver to optimize the routing schedule, given the input. The output is a new routing schedule. With heuristic B, we again allocate all the delivery units from the list to the new routing schedule. In this heuristic we use other settings determining which shifts to be more utilized than others. Then the RVRP solver is again used to optimize the routing schedule, which is the final stage of our algorithm. We have now created a routing schedule and determined how much delivery units to be delivered for each customer. These delivery quantities are based on the forecasted volumes and serve as a target volume, we can deviate from when actual volumes deviate from forecast.

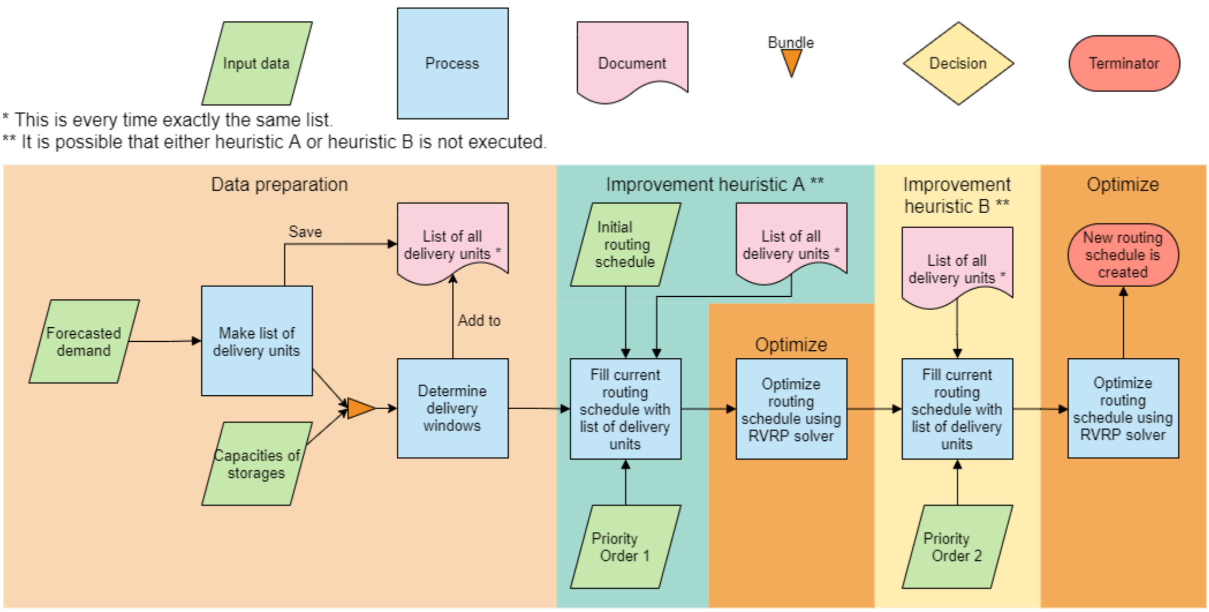

Figure 10: Process of optimizing the tactical routing schedule

Because HAVI does not know the capacities of the customers' storage locations, the data preparation phase is more extensive. In Figure 11 we show the data preparation phase when storage capacities have to be determined. Lower bounds for the capacities are first determined based on historical data. The process of creating the list of all delivery units, which is based on the forecasted demand is the same as the process shown in Figure 10. However, the delivery windows are determined in a different way. We add the assumption that the current routing schedule consists of only feasible delivery patterns. Because of this assumption we know that every delivery unit should have at least one possible delivery moment in which it can be delivered. When it is found that a delivery unit has no feasible delivery moment, we conclude that the input for our model is too restrictive, because we assumed that current routing schedule consists of only feasible delivery patterns. After adjusting the input data, we update the delivery window and check whether there is a feasible delivery moment again. After all delivery units have at least one feasible delivery moment, we continue the same process as shown in Figure 10 from heuristic A and further.

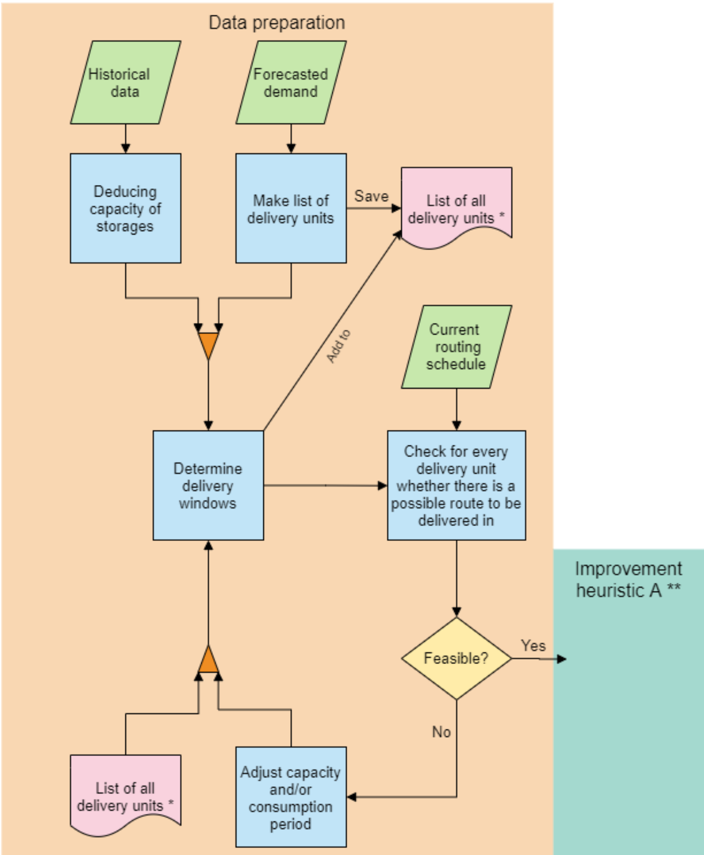

Figure 11: Data preparation including capacity determination

4.2 The unique approach of this model

We build a model in which the delivery days and quantities are determined in one phase, by determining the possible delivery moments for each delivery unit. The RVRP solver has the freedom to spread the delivery units over different delivery moments. In this research, we construct the model for HAVI, but the logic behind the model is applicable for comparable problems. Before we continue, we explain some terminology: time and delivery windows.

- Time windows: The periods of time a customer can receive deliveries (see Appendix I)
- Delivery windows: the periods in which a delivery unit can be delivered. The following factors determine the delivery windows for a specific delivery unit:
  - Weekly demand volume for the specific temperature zone.
  - Balance of demand during the week.
  - Storage capacity for the specific temperature zone.
  - Shelf life agreements with customer.
  - Time windows of the specific customer (The delivery windows are always a subset of the time windows of the corresponding customer).

To make a better comparison with the traditional approach we show differences in Table 2. Appendix VI shows an example of how switching partial deliveries can obtain results that could not be achieved with traditional approaches.

Table 2: Comparison approaches

Traditional approaches from literature	Our approach
Delivery frequency is either fixed or flexible	Delivery frequency is flexible
Delivery days and quantities are being determined in different phases	Delivery days and quantities are being determined in a single phase
Improvements by switching complete deliveries	Improvements by switching delivery units
Shelf life is sparsely applied	Shelf life agreements are included
Inventory policy should be given	Inventory limits are included, but there is no inventory policy

4.3 Input to determine customers' storage capacities

If capacities are already known, this section can be skipped. In this section we discuss which data we use to determine the capacities of the customers' storages. Before we continue, we first assume that we always deliver full roll containers. In reality, this is not always the case. The input data required for our model consists of the following:

- Historical deliveries data.
- Forecast for the near future.
- Sales distributions.

Historical deliveries data

We assume that what has been delivered in the past has fit in the storage. Therefore, we need the number of delivery units per temperature zone per delivery. HAVI uses roll containers and dollies. In a truck, three roll containers take as much space as four dollies. We assume in the rest of our research that a dolly is always 0.75 roll container.

In the system we can look back for five years. However, data on temperature zone level, can only be downloaded for the past hundred days. To increase the performance of the model, we advise HAVI to extent the data by downloading the historical data at least every hundred days. As an example, we show the first eight lines of the data download in Table 3. For every delivery there is one line per temperature zone and two lines within the frozen area, one line for the roll containers (Code 911) and one line for the dollies (Code 903). The first column represents the date of departure of the route, which in our case is the same as the delivery day. The second column represents the ID of the customer. The third column is the number of pick units. The fifth column represents the abbreviation of the business. The sixth column represents the code of the type of pick unit. The right-most column represents the temperature zone.

Table 3: Example of historical deliveries data

Departure date	Cust ID	Quantity	CUST Business ID	Unit Code	Temp Zone Description
27-01-2020	1025	1	MCD	903	Frozen
27-01-2020	1025	1	MCD	911	Chilled
27-01-2020	1025	1	MCD	911	Dry
27-01-2020	1025	2	MCD	911	Frozen
27-01-2020	1033	2	MCD	903	Frozen
27-01-2020	1033	1	MCD	911	Chilled
27-01-2020	1033	1	MCD	911	Dry
27-01-2020	1033	2	MCD	911	Frozen

Forecast for near future

The tactical routing schedule that we build will be active for a period of months. The period should be chosen such that the volumes over the weeks do not differ too much. Although the forecasting of volumes is out of scope for this research, we explain the following factors that influence the volume forecast:

1. Near past.
2. Budget.
3. Sales forecast restaurant.
4. Growth factor.

We look at the near past to see what the weighted average demand over the past three weeks is, in which the week with highest volumes gets the highest weight. We use this weighting because we want the routes to have sufficient capacity in most of the weeks. The budget also influences the forecast, because the agreements with the customer are made that they are responsible to give this input. We also use the expected forecast for the coming weeks that is filled in by every restaurant. Finally, we apply a growth factor. This factor is based on (i) seasonality, and in this specific time, (ii) on changing legislations regarding the Covid-19 measures. We take a weighted average of all these factors to conclude are forecast. The volumes for McDonald's are, despite Covid-19, at this moment already near the volumes of 2019. However, volumes per store fluctuate a bit more.

Sales distribution

The sales distribution indicates the fraction of sales per hour of the week per restaurant. This hourly demand distribution is based on historical sales data. An example of such a distribution is shown in Table 4. The data used is based on sales from July 2018 until June 2019. Every restaurant has its own hourly sales distribution. In this example we see the sales fractions from Monday 6 am till Monday 1 pm. The first hour of the table, the restaurant is closed, or sales per hour is lower than 0,05% of weekly volume. The hours after that, sales increase every hour.

Table 4: Example of hourly sales distribution

Day	Mon	Mon	Mon	Mon	Mon
t (hour)	5	6	7	8	9
From	8:00 am	9:00 am	10:00 am	11:00 am	12:00 pm
To	9:00 am	10:00 am	11:00 am	12:00 pm	1:00 pm
% of weekly sales	0,0%	0,1%	0,3%	0,4%	1,1%
Cumulative	0,0%	0,1%	0,4%	0,8%	1,9%

4.4 Deducing capacity of storage locations

If capacities are already known, this section can be skipped. In the previous section we discussed the input needed to deduce the capacity of storage locations. In this section we are going to determine lower bounds for the capacities.

In earlier simulation studies HAVI used the surface of the storage as indicator of the storage capacity. However, this is not accurate. For example, some freezers are only accessible via the fridge, therefore, more space in the fridge must be reserved as a walking path. This is just one out of many examples why the surface of the storage location is a poor indicator for the capacity. In this research, we deduce the storage capacity based on historical data as well as on expected forecasted volumes. We define two criteria for which we assume that they are true in practice. Based on these assumptions we deduce lower bounds for the storage capacities. The two criteria are as follows:

1. The capacity is at least as big as the largest delivery that is received in the past.
2. The capacity is at least as big as the largest amount that is expected to be consumed between two deliveries within the planning horizon.

Using these criteria, we know the lower bound of the capacity. In practice, the capacity can be higher. Therefore, we do not utilize the full potential. Further actions or research needs to be done to increase the knowledge of the storage locations' capacity.

Criterium 1

We assume that the obtained historical data of the deliveries is correct. This leads to knowing the size of the largest delivery that has taken place in history. In this criterium we assume that the capacity of the storage location is at least enough to have received the largest delivery.

One reason why the capacity is not at least the size of the largest delivery, is that the delivery in the past maybe did not fit in the storage. This could be when delivery took place before a huge sales peak, and not all goods had to be placed in storage, but a part was directly placed in the kitchen to be used. On the contrary, at arrival of a delivery, there are always still some products in storage. Therefore, the capacity is bigger than what is added to storage. To combine the delivery data with the data stating how much goods were already in storage upon delivery, is left for further research. Furthermore, because we have data of the past hundred days, the period we are looking at are either winter months or months in which Covid-19 had its impact. Therefore, we take the largest delivery from this period.

In the future, HAVI has to think of a method how to exclude the outliers. For example, on king's day, deliveries can grow so large, that it is uncomfortable for the restaurants but accepted as an exception.

Criterium 2

In this criterium, we assume that all consumption between two deliveries (i.e., delivery A and B) fits in storage. All goods that are being consumed between delivery A and B, were delivered in delivery A or earlier. Because they are in storage just after delivery A, but not yet consumed, they were in the storage all at the same time.

The exception to this logic is when a restaurant also has received goods between delivery A and B that were not delivered by HAVI. Since HAVI is the only external supplier, the only option is that they received goods from other neighboring restaurants (i.e., mutual delivery). If a restaurant did receive a mutual delivery between delivery A and B, it is not true that the capacity of the storage is at least as big as the amount that is consumed between delivery those deliveries.

However, we do assume that the volume consumed between two deliveries fits in the storage. If two deliveries are planned such that the sales forecast between the two deliveries is greater than the capacity of the storage, the restaurant planner must adjust the order such that it fits in the storage. The delivery is then smaller than needed to cover all sales to the next delivery. Therefore, a mutual delivery is needed. The assumption that the volume being consumed between two deliveries fits in the storage, can lead to the restaurant needing a mutual delivery, in the case that they in practice also needed a mutual delivery.

In this research, we use the timing of deliveries in the current situation in combination with the forecasted volumes and historical sales distribution to determine the maximum volume that is being consumed between two deliveries. For example, based on the data in Table 4, the maximum amount between two deliveries is 40% of the weekly volume. We assume that 40% of the expected weekly volume fits in the storage. For this statement we use the fact that in the current situation the delivery patterns are determined or approved by the restaurant planners. Therefore, we assume that the delivery patterns in the current situation are feasible, given the forecasted volumes. From this assumption it follows that the largest expected volume consumed between two deliveries, fits in the storage. We do not take the actual maximum amount being consumed between two deliveries, because the sales data that is available is expressed in euros. This cannot be translated to storage space. To determine the corresponding storage space that is being consumed between two deliveries, we need to know for each article how often it is sold and how much storage space is used for that article. We do not have data on this detail level.

Summarizing

We defined two criteria and assume them to be true in practice. Therefore, the storage capacities in practice are at least as big as the criteria are stating. We therefore have determined lower bounds for the capacities. In practice, the storages may be larger.

4.5 Determining delivery windows

The basis of our approach is to determine the delivery window for every delivery unit. A delivery windows is subset of the time windows corresponding to the specific customer bounded by an earliest and latest delivery moment. The earliest and latest delivery moment possible for delivery unit U will be denoted by $Earliest(U)$ and $Latest(U)$, respectively.

Earliest possible delivery moment

The earliest possible moment of delivery is affected by the following aspects:

- Period of consumption.
- Shelf life agreements.
- Storage capacity for specific temperature zone.
- Possible delivery moments of previous delivered delivery units.

We distinguish the limits by two categories. The first category contains the first two aspects and are related to the shelf life. The other category contains the last two aspects and are related to the storage capacity. Before elaborating more on these categories, we look at the consumption period of a delivery unit.

Consumption period

The consumption period of a delivery unit is the time that is used to consume the goods on a delivery unit. Without losing generality, we assume that consumption periods of delivery units are strictly disjunctive (i.e., do not overlap). Furthermore, we assume that the start of the consumption period of the first delivery unit in a week, is exactly at the beginning of the week (i.e., Monday 4 am).

- The weekly demand for customer C for temperature zone T is denoted by $D_{C,T}$.
- The n[th] delivery unit of customer C for temperature zone T is denoted by $U_{C,T,n}$.
- The start of the consumption of $U_{C,T,n}$ is when $n - 1/_{D_{C,T}}$ part of weekly volume is consumed, and this is denoted by $t_{C,T,n}^{start}$.
- The end of the consumption of $U_{C,T,n}$ is when $n/_{D_{C,T}}$ part of weekly volume is consumed, and this is denoted by $t_{C,T,n}^{end}$.

$t_{C,T,n}^{start}$ and $t_{C,T,n}^{start}$ are both being expressed in hours, starting with Monday 4 am = 0, Monday 5 am = 1 until Monday 3 am = 167. For example, when $D_{C,T} = 4$, the start of the consumption period of the 2^{nd} delivery unit for customer C in temperature zone T is when $2 - 1/_{4} = 25\%$ of weekly volume is consumed. From the cumulative data as in Table 4, we know which hour corresponds to that percentage.

Shelf Life

For frozen and dry goods, the shelf life is longer than a week. Since we assume that we do not deliver goods more than a week earlier, the minimum shelf life agreement (SLA) for frozen and dry goods is set at seven days. The minimum shelf life agreement for chilled goods is four days, including the day of delivery (e.g., Monday delivery until Thursday consumption). In our study we define the start of a day as 4 am. Every time between midnight and 4 am will belong to the previous day. We set the earliest delivery moment possible, regarding the shelf life, based on the end of the consumption period. If the end of the consumption period for the n[th] delivery unit of customer C for temperature zone T is during day d, the earliest delivery moment, regarding the shelf life, is at the start of day $d - (SLA - 1)$ and expressed in Equation 2. $t_{C,T,n}^{end}$ is expressed in hours. To know at which day this hour is, we divide by 24 and round down. We subtract three days and multiply by 24 again, to get the hour representing the start of the day, $SLA - 1$ days before the end of the consumption period.

Equation 2

$$Earliest(U_{C,T,n})_{Shelf} = (\lfloor t_{C,T,n}^{end}/24 \rfloor - SLA) * 24$$

Capacity

The second category limiting the earliest delivery moment possible is related to the storage capacity.

- The capacity of the storage location for temperature zone T for customer C is denoted by $Cap_{C,T}$.
- We assume that the delivery units are being delivered in the same order as in which they are being consumed.
- Capacity should be enough also in a scenario where all volumes are equal to 110% of forecasted volume.

We take 110% of forecasted volume because it is reasonable that the actual volumes will deviate less than 10% of forecasted volume. On average the deviation for the total market sales per day is about 4%. The safety of 10% is on customer level per delivery. A delivery contains possibly multiple days of sales. Based on correlation factors of sales of consecutive days and including a weighting how many days are between two consecutive deliveries, the percentage could be determined, for which holds true that 95% of all deviations are smaller than this percentage. However, we have taken this 10% as a rule of thumb, which is reasonable looking at the numbers.

Because of the assumption of the order of delivery, we know that upon delivery of delivery unit $U_{C,T,n}$ at time t all delivery units $\{U_{C,T,i} \mid i < n \wedge t_{C,T,i}^{start} > t\}$ are in storage (i.e., at the moment of delivery, every delivery unit that is already delivered, but not yet consumed, is still in storage). Here, t is also expressed in hours. We state that the earliest delivery moment regarding the capacity restriction is:

Equation 3

$$Earliest(U_{C,T,n})_{Cap} = t_{C,T,(n-Cap_{C,T})/1.1)}^{end}$$

To deliver a delivery unit, there should be a place available in storage. $U_{C,T,(n-(Cap_C^T)/1.1)}$ is the delivery unit $Cap_{C,T}$ places before delivery unit $U_{C,T,n}$ in the 110% volume scenario. In Equation 3 we say that the earliest possible delivery moment of $U_{C,T,n}$, regarding the capacity, is equal to the end of the consumption period of $U_{C,T,(n-\frac{Cap_{C,T}}{1.1})}$. An example with capacity set at two, determining the delivery window for the second delivery unit of four delivery units per week is as follows:

$Earliest(U_{C,T,2})_{Cap} = t_{C,T,0.18}^{end}$ is when ${}^{0.18}/_{4} = 4.5\%$ of weekly volume is being consumed.

The shelf life and the capacity both have a limitation on the earliest delivery moment possible. Therefore, we conclude from Equation 2 and Equation 3:

Equation 4

$$Earliest(U_{C,T,n}) = max(Earliest(U_{C,T,n})_{Shelf}\ , Earliest(U_{C,T,n})_{Cap})$$

Latest possible delivery moment

The latest possible delivery moment is easier than the earliest possible delivery moment. The latest possible delivery moment is limited by the consumption period. It is trivial; the delivery must take place before consumption starts, concluding:

Equation 5

$$Latest(U_{C,T,n}) = \ t_{C,T,n}^{start}$$

Appendix VII contains an example of how delivery windows are determined. This example also discusses how to deal with the following:

- Crossing the borders of a week.
- Periods in which no volume is consumed.

4.6 Further relaxation of lower bounds

If capacities were already known, this section can be skipped. Also, when is chosen that a routing schedule is constructed from scratch instead of using an initial solution, this section can be skipped. In Section 4.4, we stated two criteria, which determined lower bounds for the capacity of the storages. In this section we show that if we assume that the current routing schedule consists of feasible delivery patterns, we can further relax those lower bounds.

4.6.1 Infeasible situations

The delivery patterns in the current situation are determined or approved by the restaurant planners. We therefore assume that the delivery patterns are feasible in practice. However, it is possible that they are not yet feasible in our model. There are two situations why the current routing schedule according to the input data is not feasible:

1. The capacity we deduced is too restrictive.
2. The consumption period we determined is too restrictive.

Situation 1

It is possible that the capacity that we discussed in Section 4.4 is not sufficient. There are two reasons why this could be the case:

a. The consumption periods are not in sync with the delivery moments.
b. The delivery window is smaller than the time between two deliveries

Reason a

We give an example of a delivery window that lies between two consecutive delivery moments. Therefore, the delivery unit cannot be delivered in any of the two delivery moments. Suppose that in Section 4.4 we deduced the capacity of a specific storage to be $Cap_{C,T}$, which is the maximum amount being consumed between two deliveries. Suppose that the delivery moments in the current schedule are such that the first delivery (delivery A) takes place after 0.5 delivery units are being consumed and the second delivery (delivery B) takes place after $Cap_{C,T} + 0.5$ delivery units are being consumed. We will show that the delivery window of $U_{C,T,(Cap_{C,T}+1)}$, which is $U_{C,T,n}$ with $n = Cap_{C,T} + 1$ lies between those two boundaries. From Equation 3 it follows that

$Earliest(U_{C,T,(Cap_{C,T}+1)})_{Cap} = t^{end}_{C,T,((Cap_{C,T}+1)-(Cap_{C,T})/1.1)} > t^{end}_{C,T,1}$. Therefore, the start of the delivery window is later than delivery A. From Equation 5 it follows that

$Latest\left(U_{C,T,(Cap_{C,T}+1)}\right) = t^{start}_{C,T,(Cap_{C,T}+1)} = t^{end}_{C,T,(Cap_{C,T})}$. Therefore, the end of the delivery window is earlier than delivery B. Therefore, this delivery unit cannot be allocated to any of the routes in the current situation and therefore, no feasible delivery moment exists.

Reason b

Suppose that the result of criterium 2 is that the capacity of a specific storage is $Cap_{C,T}$, because that is the maximum expected amount to be consumed between two deliveries (i.e., there are $Cap_{C,T}$ consumption periods between delivery A and B). When we determine the delivery window, the earliest delivery moment is determined in Equation 3 and set equal to $t^{end}_{C,T,(n-(Cap_{C,T})/1.1)}$. The latest delivery moment is determined in Equation 5 and set equal to $t^{start}_{C,T,n}$, which also equals $t^{end}_{C,T,(n-1)}$. Therefore, the width of the delivery window is $(n-1) - \left(n - \frac{Cap_{C,T}}{1.1}\right) = \frac{Cap_{C,T}}{1.1} - 1$ consumption periods, which is smaller than the $Cap_{C,T}$ consumption periods between delivery A and B. Both reasons explained how it can be that the capacity that we determined is too restrictive. Apparently, the capacity in practice is larger than we stated in the model. When this is the case, we increase the value of the capacity with 10%. How this is applied is explained in Section 4.6.2.

Situation 2

We assumed that all delivery units are full (i.e., we assumed every delivery unit to contain products covering the same amount of sales). However, in practice the utilization of delivery units can differ. When a delivery unit is not fully utilized, its consumption period covers less sales. Especially in the restaurants with small storages, restaurant planners do not utilize the delivery units fully. Suppose that three deliveries are spread such that 1.2 roll container is consumed between delivery A and B, and 0.9 delivery unit between delivery B and C. When, the storage is not much larger than 1.2 roll containers, the restaurant planner will not make the order larger than the 1.2 roll container that is strictly needed. The first delivery consists of two delivery units that are 60% utilized on average. The next delivery consists of one delivery unit, that is 90% utilized. The consumption period of the last delivery unit contains 50% more sales than the first two. Because we assumed only full delivery units are being delivered, we also assumed the consumption periods to contain equal amounts of sales.

We see that the sales per delivery unit can differ in practice, but we have assumed that every delivery unit represent equal amount of sales. Therefore, it is possible that some delivery units can be delivered in practice, but not in our model. To adjust for this effect, we shorten the consumption periods for these delivery units. Because the delivery window is dependent on the start of its own consumption period, we shorten the consumption period by adjusting the start of the consumption period (i.e., setting it to a later moment). When the start of the consumption period is at hour h and the consumption period's duration is d, we adjust the start of the consumption period to $h + 0.1 * d$. The consumption period is now decreased with 10%. This 10% is chosen arbitrarily. We know that the consumption period is too long, but not how much too long. This 10% is used in the algorithm in the next section. This algorithm now runs in about two minutes. Decreasing the 10% makes the algorithm more accurate at the expense of calculation time.

4.6.2 Algorithm feasibility

In Section 4.6.1 we gave some situations in which it is possible that due to the input being too restrictive, some delivery units cannot be delivered in any of the delivery moments, although the delivery pattern is feasible in practice. For all delivery patterns from the current routing schedule to become feasible in our model, all delivery units should have at least one possible delivery moment in the current routing schedule. To do this, we need an algorithm to adjust the input data (i.e., capacities and shelf life agreements that are too restrictive). We show the process diagram of the algorithm in Figure 12. Using the list with all delivery units, we loop over all customers and all temperature zones, to check whether there is a delivery unit without a feasible delivery moment in the current routing schedule. If there is a delivery unit without a feasible delivery moment, we check whether it is due to the capacity constraints, or the shelf life constraint.

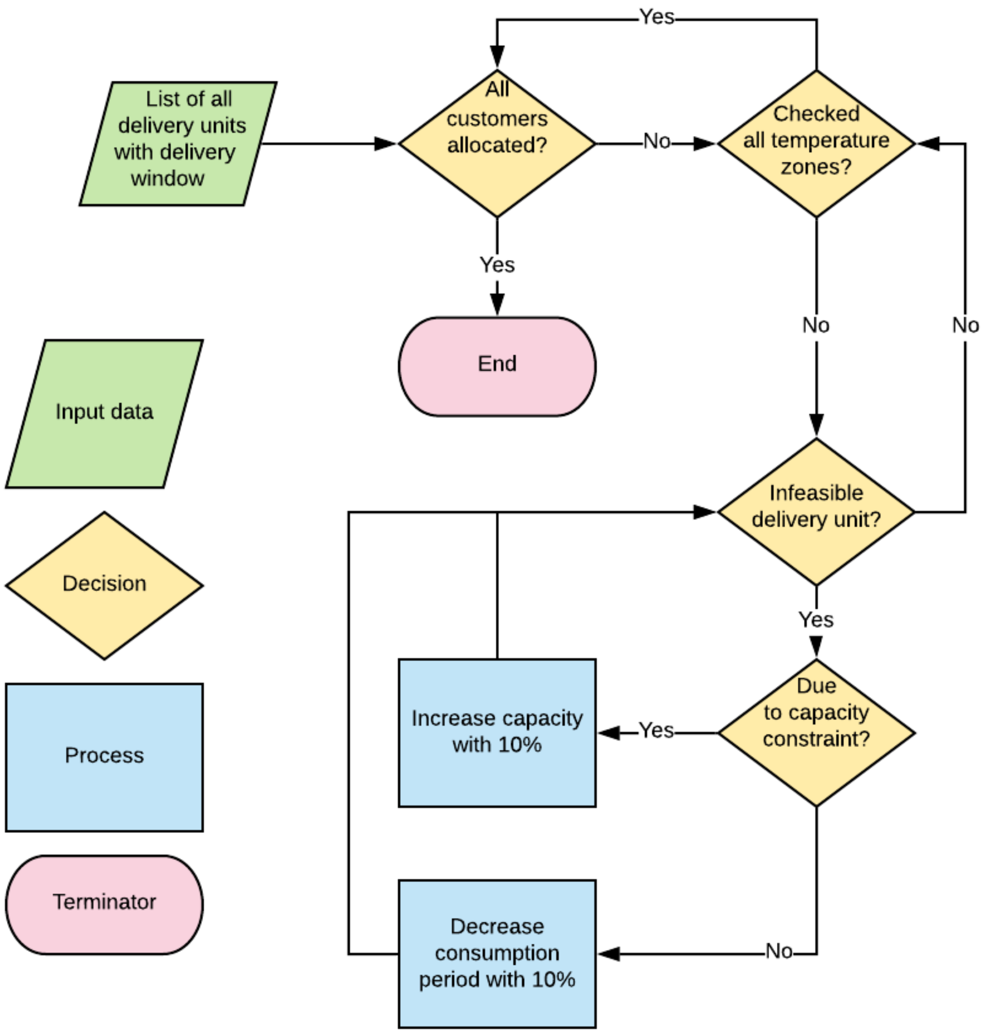

Figure 12: Process Diagram Algorithm Feasibility

From Equation 4 it follows that at least $Earliest(U_{C,T,n}) = Earliest(U_{C,T,n})_{Shelf}$ or $Earliest(U_{C,T,n}) = Earliest(U_{C,T,n})_{Cap}$. If the capacity of the customers storage is the limiting factor, we know that the derived capacity is lower than the actual capacity. Therefore, we increase the input value of the capacity with 10%. Otherwise, the shelf life is the limiting factor and we reduce the consumption period with 10% by adjusting the start of the consumption period and setting it to a later moment.

4.7 Algorithm for optimization

We now develop our main algorithm. We are done with the data preparation phase and continue with the improvement heuristics. In Figure 10, we see how this relates to the overall process. Improvement heuristic A and B are used to escape from local optima. Heuristic A emphasizes the removal of unnecessary deliveries in the most costly or largest (based on workload) shifts and adding deliveries in the less costly and smaller shifts. Heuristic B emphasizes the utilization of the least costly and smallest shifts. In Chapter 5, we experiment with applying either one or both the improvement heuristics.

4.7.1 Tradeoff costly shifts vs workload balancing

In this section, the terminology ‘priorities of shifts’ is important to understand. The priority of the shifts is defined such that the shifts with the lowest costs and/or lowest workload have the highest priorities. We want to utilize the shifts having a high priority more than in the current situation. Several times in our algorithm an allocation procedure will be used. The allocation procedure emphasizes either on the shifts with the lowest priorities or the shifts with the highest priorities.

Some delivery units have a small delivery window and must be delivered on a specific day. Other delivery units have a broader delivery window and in the allocation procedure we determine on which route they are delivered. In Figure 13 we see that an emphasis on the shifts with the lowest/highest priorities results in routes on the most/least expensive shifts to be utilized most. In different stages of the algorithm, the focus is either on the low or high priorities. These allocation procedures are used to escape from local optima. Afterwards the RVRP solver optimizes the routing schedule.

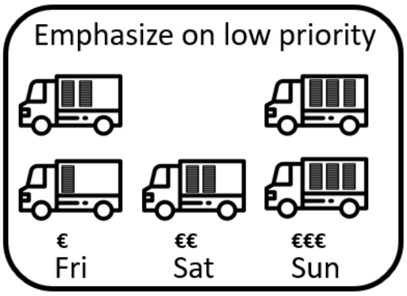

Figure 13a: Emphasize on low priority

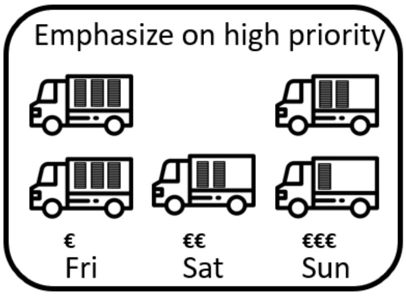

Figure 13b: Emphasize on high priority

Giving priorities to the shifts we consider two factors:

- Costs related to a shift (in the weekends salary supplements are in place).
- Workload of a shift.

Costs

The costs of a tactical routing schedule consist of the following:

- Labor costs.
- Marginal costs for distance.
- Costs for the vehicles.

Labor costs

Labor costs are split into costs related to transport and costs related to warehouse. Labor costs related to transport exists of the duty time of a driver. The duty time for a route is measured from departure from the DC, until arrival at the DC, increased with one and a half hour. The increment of one and a half hour is used because this time is needed for loading, unloading and the administration of a route. We set the hourly costs at €30. Labor costs related to the warehouse are determined based on the number of delivery units that must be picked. The total costs of all handlings are set at €4 per delivery unit. For both transportation and warehousing, salary supplements are included as stated in Table 5. The weighted average is based on the approximation that one third of the costs is warehouse related and two third is transport related.

Table 5: Costs factor per shift in the weekend

Delivery moment	Pick moment	Weighted average
Saturday daytime (150%)	Friday evening (100%)	133%
Saturday evening (150%)	Saturday daytime (150%)	150%
Sunday daytime (200%)	Saturday evening (150%)	183%
Sunday evening (200%)	Sunday daytime (200%)	200%
Monday daytime (100%)	Sunday evening (200%)	133%

Marginal costs for distance
The main components of the marginal costs for distance traveled are (i) the fuel used and (ii) the maintenance costs that are allocated to mileage. These costs together are set at €1 per kilometer.

Costs for the vehicles

The main components for the costs related to the fleet size are the costs for leasing the trucks, insurance, and maintenance for damage that is not covered and allocated to aging. These costs together are set at €2000 per week per truck. In our results we measure the deviation in costs with respect to the current situation (e.g., 22 trucks have zero costs).

Summarizing, the costs used for this research are as follows:

- €1 per kilometer traveled.
- €30 per hour for a driver.
- €4 per delivery unit handled in the warehouse.
- €2000 per truck per week.

Workload balancing

Besides reducing costs based on supplementary salary, we also aim at balancing the workload. In Figure 14 we show the number of routes that are planned per shift. When focusing on workload balancing, we prioritize based on these numbers. When two shifts have the same score, the prioritization is based on secondary arguments (e.g., a day shift is often more efficient to schedule). Workload balancing is scored based on penalizing deviations from the preferred workload per shift.

Subnetwork

Because in this research we only consider data from a fifth part of the total network, we need to measure the workload balancing in how it contributes to the balance of the total network. If this study would be applied to the total network, the part until Table 6 doesn't have to be executed, but it is good to read for terminology. The workload balance for the selection that is within the scope of this research is not the same as the workload balance in the total network. The preferred workload (delivered delivery units) per shift is determined such that it contributes to balancing the workload in the total network.

- The current workload of shift s in the Netherlands is denoted by W_s^{NL}.
- The current workload of shift s in the east of the Netherlands is denoted by W_s^{East}.
- The target workload of shift s in the Netherlands is denoted by T_s^{NL}.
- The target workload of shift s in the east of the Netherlands is denoted by T_s^{East}.

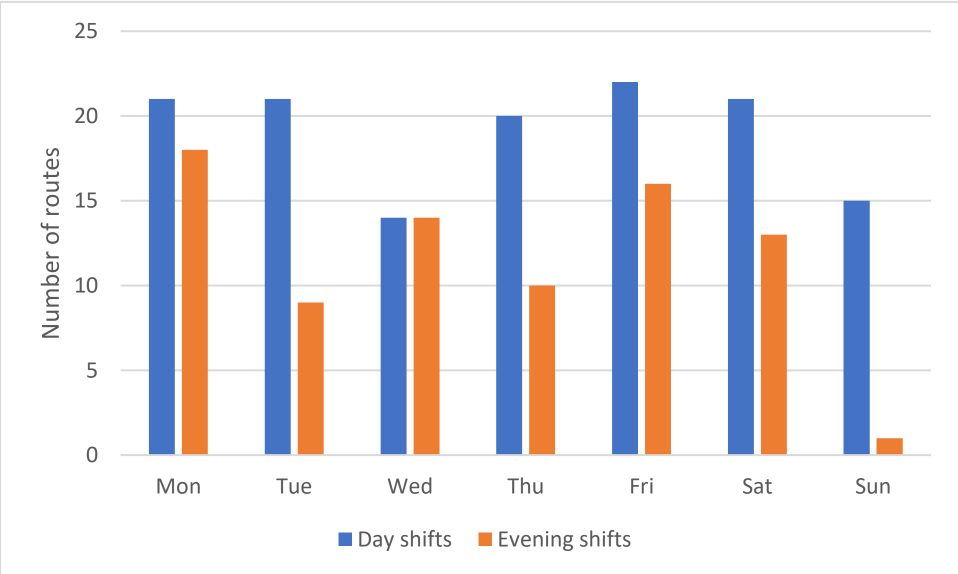

Figure 14: Routes per shift

We determine the target workload for shift s in the east of the Netherlands as follows:

Equation 6

$$T_s^{East} = W_s^{East} - (W_s^{NL} - T_s^{NL})$$

The results of Equation 6 for each shift are stated in Table 6 (D = Day, E = Evening). We now have determined targets for the workload for each shift. If we used data of the total network instead of only a part (as we do in this research with the eastern part of the Netherlands), we could just have filled in the preferred workload balance for the total network. In the remainder of this calculations this preferred values have to be filled into the variables T_s^{East} although it applies for the total network.

Table 6: Target workload East of the Netherlands

Shifts	Mon D	Mon E	Tue D	Tue E	Wed D	Wed E	Thu D	Thu E	Fri D	Fri E	Sat D	Sat E	Sun D	Sun E
Current NL	1273	900	1205	457	1001	903	1197	663	1260	968	1273	798	988	9
Current East	183	284	164	117	162	293	175	137	151	279	192	229	201	0
Target NL	991	991	991	991	991	991	991	991	991	991	991	991	991	9
Target East	-99	375	-50	651	152	381	-31	465	-118	302	-90	422	204	0

The workload of shift s in experiment e is denoted by $W_{s,e}$. The total penalty of the workload distribution (WP) of experiment e is calculated as follows:

Equation 7

$$WP_e = \sqrt{\sum_{s=1\ to\ 14}\left(W_{s,e} - T_s^{East}\right)^2 / 14}$$

The delta between $W_{s,e}$ and T_s^{East} is being squared because the larger the difference, the larger the marginal extra costs we assume. We divide this by 14 and take the square root to normalize the value of the penalty. It is not equal to the average deviation because the higher the deviation, the more severe the penalty due to the square. When WP_e equals zero it means that the workload in each shift is exactly equal to the target level.

4.7.2 Improving tactical routing schedule

When is chosen that a routing schedule is constructed from scratch instead of using an initial solution, this subsection can be skipped. We now explain the algorithm to improve the tactical routing schedule. Remember that this is not a construction heuristic and we thus need an initial solution. For HAVI, we use the current routing schedule as initial solution. The algorithm is split in improvement heuristics A and B. Both based on the following allocation procedure. This allocation procedure is needed to escape from local optima while steering on which shifts to utilize most.

Allocation procedure

The allocation procedure is used to determine to which routes the delivery units are being allocated, as is visualized in Figure 13. The priority of the shifts is already determined. In Section 4.6, we ensured that each delivery unit can be assigned to at least one possible route (i.e., delivery moment). The delivery units that can be delivered in only one route must be allocated first, because they have no other option. For delivery units that can be delivered to multiple routes, we must prioritize which delivery units to allocate first. The delivery units with the least number of options, and the delivery units for which the alternative options are more expensive should be routed earlier than delivery units with multiple inexpensive options. The To prioritize this way, we build the following allocation procedure:

(i) Assign points to the shifts such that shifts that needs to be utilized most, have the highest number of points. An example is presented in Table 7, here we focus on reducing costs first and balancing workload second. The priorities follow from Table 5 and Figure 14. In this example, the higher the priority, the higher the number of points. Therefore, we emphasize on high priority as in Figure 13b. All other combinations to assign points are presented in Appendix VIII.

Table 7: Points assigned to priorities of shifts

Priority	#1	#2	#3	#4	#5	#6	#7	#8	#9	#10	#11	#12	#13	#14
Shift	Tue E	Thu E	Wed D	Wed E	Fri E	Mon E	Thu D	Fri D	Tue D	Mon D	Sat D	Sat E	Sun D	Sun E
Points	100	90	80	70	60	50	40	15	12	7	4	3	2	1

(ii) For all delivery units, determine to which routes and corresponding shifts it can be allocated to.

(iii) Assign points to all delivery units. The points are being equal to the summation of all shifts the delivery unit can be allocated to. For example, a delivery unit with a delivery window from Tuesday 8 am until Wednesday 6 pm, where the customer has only a delivery (in the current routing schedule) in a route on Wednesday daytime, gets 80 points.

(iv) Loop over all routes $r = 1\ to\ R$ starting with the routes in the shifts with the highest points. r is the counter, and R is the total number of routes.

(v) For all delivery units that can be assigned to route r, determine how many delivery units can be delivered to route r, have less or equal number of points. Delivery units that are already allocated to another route are excluded.

(vi) Determine the number of available places (AP_r) in the route for McDonald's delivery units. The delivery quantities of other customers are being fixed, as assumed in Section 1.6.

(vii) Number the delivery units from lowest to highest points, all delivery units get a unique number nr . Therefore, delivery unit nr is the nr^{th} delivery unit that will be allocated.

(viii) Delivery unit nr is allocated to route r if $nr \leq AP_r$.

With this allocation procedure it is possible that there are delivery moments to which no delivery units are being allocated. Likely, the delivery moments not being used are in shifts we do not emphasize. If the emphasis was on high priority shifts, we remove the delivery moments that are not being used from the current routing schedule. However, if the emphasis was on low priority shifts, we do not remove those delivery moments, but add a dummy delivery unit to this delivery moment. This dummy delivery unit is explained in Section 4.7.3 at step 5.

We now explain heuristic A and B. It is possible to either execute only one of the heuristics without the other, or to first execute heuristic A followed by heuristic B. If heuristic B is executed after heuristic A, then the result of heuristic A will be input for heuristic B. If we do not first execute A, then the current routing schedule is the input for heuristic B.

Heuristic A

The goals of heuristic A are:

(i) Avoid unnecessary deliveries in the most costly or largest shifts (i.e., shifts with the lowest priorities).

(ii) Add deliveries to the shifts that are less costly or which workload is lower (shifts with highest priorities).

We want to avoid that we use a delivery moment, just because it is in the current routing schedule. Especially delivery moments in shifts with lowest priority we want to avoid. Therefore, we start filling routes in the shifts with the highest priorities as much as possible. Only when some delivery units cannot be allocated to the shifts with the higher priorities, we must deliver it in the shift with the lower priority. When no delivery units are allocated to a certain delivery moment, apparently this delivery moment is not needed to meet customers' storage restrictions. These deliveries are removed from the current routing schedule. All routes are emptied, and again all delivery units are being allocated, but this time with emphasis on low priority shifts. The advantage of emphasizing on the low priority shifts is that the routes in these shifts have spare capacity such that new delivery moments can be added in these routes. To prevent the RVRP solver to merge routes, we add a dummy customer in each route. These dummy customers are defined such that they cannot be delivered in the same route.

Heuristic B

The purpose of heuristic B is to move the workload as much as possible to shifts with the highest priorities. In the allocation method the delivery units are assigned as much as possible to the shifts with the highest priorities. The goals of heuristic B are:

(i) Use the opportunities to merge routes in the shifts with the lowest priorities because they have a low utilization. This way, we reduce the number of routes in those shifts.

(ii) Increasing the utilization of the shifts that (i) have a low workload or (ii) are less costly.

The steps of the algorithm

The Algorithm we use to improve the current tactical routing schedule consists of nine steps. The first seven steps belong to heuristic A and the last two steps belong to heuristic B. The algorithm contains the following steps:

Heuristic A

1. Set the order of the shifts such that the highest priority is given to the shift that we want to utilize the most.
2. Use the allocation procedure with emphasis on high priority.
3. Remove the unused of the current routing schedule.
4. Use the allocation procedure with emphasis on low priority.
5. Add a dummy delivery unit to every delivery to which no delivery unit is assigned.
6. Add a dummy customer to each route, preventing the RVRP solver to merge routes.
7. Using the RVRP solver, optimize the routing schedule based on duty time.

Heuristic B

In some experiments we already applied heuristic A. The outcome of step 7 (without the dummy customers) serves as input for heuristic B. Otherwise, we take the current routing schedule as input.

8. Use the allocation procedure with emphasis on high priority.
9. Using the RVRP solver, optimize the routing schedule based on duty time.

4.7.3 The algorithm elaborated

When is chosen that a routing schedule is constructed from scratch instead of using an initial solution, this subsection can be skipped. In the previous section we mentioned the steps of our algorithm on how to optimize our routing schedule. In this section we elaborate more on some step.

Step 1 and 2

In these steps we set the priorities of the shifts and assigned points as in the example of Table 7. In this research, we are using four different scenarios, which are stated in Appendix VIII.

Step 5

In step 4 the emphasis was on filling the routes within the shifts with low priority. Possibly, to some delivery moment (especially in shifts with higher priority) no delivery units are being allocated. Since it is not the goal to reduce the number of deliveries in the shifts with higher priority, we add a dummy delivery unit to these deliveries. We decide to set the delivery window of the dummy delivery unit equal to the adjacent days in such a way that no other delivery moment from the current routing schedule can be chosen for this delivery unit. In Table 8 we show the delivery window for a dummy delivery unit that corresponds to a delivery on Tuesday, given the current delivery pattern. We see that the dummy delivery unit can be delivered anytime from Monday until Thursday.

Table 8: Delivery windows Dummy delivery unit

Day	Mon	Tue	Wed	Thu	Fri	Sat	Sun
Current pattern		x			x		X
Dummy delivery windows	----------	----------	----------	----------			

Step 6

In this step we add a delivery to a dummy customer to every route. We do not deliver any delivery units, so this customer does not impact the utilization of the truck. The dummy customer is located near to the DC and is visited first in the route. The loading time for the dummy customer is one minute, which makes sure it does not impact the routing options. We set a rule in the RVRP solver that two delivery units of the dummy customer can never be delivered in the same route. This way, the delivery units remain separated while optimizing and no routes are merged.

We must prevent that routes are transferred to shifts with a lower priority. Therefore, the delivery window of the dummy customer is set equal to the possible departure times of all the shifts that have a priority at least as high as the shift in which the route is currently driven. The possible departure times are taken because the customer is close to the DC and is visited first. The possible departure times are between 4 am and 8:30 am for day shifts and between 1 pm and 6 pm for evening shifts.

Step 9

Since we left out the dummy customers, the RVRP solver may now merge routes if necessary.

4.7.4 Interventions

We define multiple interventions for the algorithm we constructed. We first state them all five, we then give some introducing statements that are used for all the interventions. Afterwards we explain the interventions one by one. The application of the interventions and the numbers used in this research are discussed in Appendix IX.

(i) Execute either heuristic A, B or both.
(ii) Prevent the balance of number of routes per shift to get worse.
(iii) Force the RVRP solver to improve the balance based on number of routes per shift.
(iv) Force the RVRP solver to reduce the fleet size.
(v) Allow consecutive deliveries to be within 18 hours.

For all interventions, we use the following:

- In Paragon, we limit the number of vehicles because we cannot limit the number of routes. A vehicle may drive only one route per shift, so the vehicle and route are used interchangeably.
- Limit the number of vehicles available to V^{max}, which is the maximum used vehicles at the same time. In our case six vehicles are used, because at Monday and Wednesday evening we need six trucks at the same time.
- For shift s currently V_s^{cur} vehicles are used.
- The limit of number of vehicles used in shift s is set at V_s^{max}.

Intervention (i)

Experiment with executing only heuristic A, which is steps 1-7; only heuristic B, which is steps 8 and 9, and heuristic A and B, which are all nine steps.

Intervention (ii)

In some experiments we prevent the balance of the number of routes over the shifts getting worse. We do this by creating deliveries to the dummy customer:

- Make $V^{max} - V_s^{max}$ deliveries to this customer in shift s.
- These deliveries should be visited as only customer in the route.
- The delivery windows of these dummy deliveries should be such that they can only be delivered in shifts with lower priority.

Intervention (iii)

By forcing the RVRP solver to improve the workload balance, at least one route must change from shift to a shift with a higher priority. We do this as follows:

- Decide from which set of shifts a route must be transferred to a shift with higher priority.
- Set $V_s^{max} = V_s^{cur} + i$ for shifts s for which the number of routes can increase with at most i.
- i must be equal to zero for the set of shifts with the lowest priority to ensure an improvement.
- Create deliveries to the dummy customer the same way as in intervention (ii).
- Remove deliveries from the set of shifts with the lowest priorities until it is possible to save a route within these shifts with the RVRP solver.
- Add an extra delivery to the dummy customer in shift s with priority p in which the RVRP solver saved a route.
- Remove a delivery to the dummy customer in the shift with priority $p - 1$.
- Reoptimize the routing schedule with the RVRP solver by routing all deliveries that were removed earlier.

Intervention (iv)

Because we deal with a subnetwork, the fleet size is not determined by the maximum number of vehicles used in this subnetwork, but the fleet size is determined based on the impact this sub network has on the larger network. Here we explain the logic applied when we deal with the total network. In Appendix IX we explain how we deal with the sub network in this research.

- Set $V_s^{max} = V^{max} - 1$ for all shifts s for which $V_s^{cur} < V^{max}$.
- Set $V_s^{max} = V^{max}$ for all shifts s for which $V_s^{cur} = V^{max}$.
- Remove deliveries from the set of shifts for which $V_s^{cur} = V_{max}$ until it is possible to save a route in each of these shifts with the RVRP solver.

- Set $V_s^{max} = V^{max} - 1$ for all shifts in which a route is saved.
- Reoptimize the routing schedule with the RVRP solver by routing all deliveries that were removed earlier.

Intervention (v)

When creating tactical routing schedules, HAVI keeps at least 18 hours between two consecutive deliveries to a specific customer. This is a setting within Paragon. With this intervention we switch this setting off such that consecutive deliveries can take place within 18 hours of each other. In this research, this option is even more interesting than in the current situation because we split deliveries on the level of delivery units.

4.8 Verification and validation

Before going to the operational policies, we verify and validate our model for the tactical routing schedule. With verification we check whether the model in the RVRP solver (Paragon in our case) satisfies all mentioned restrictions. With validation we check whether the routing schedule in Paragon represents reality accurately.

Verification

Verifying the model, we check whether what we see is as we expected it to be. To do this, in all spreadsheets we use, we add a unique identifier to every delivery unit. So, we can always use backtracking through the spreadsheets to see which input resulted in a specific output. After a test experiment, we made an overview of all the delivery patterns that were being changed with the optimization as is shown in Figure 15.

	Current							Experiment						
Custom	Mon	Tue	Wed	Thu	Fri	Sat	Sun	Mon	Tue	Wed	Thu	Fri	Sat	Sun
1009	x		x	x		x		x		x	x		x	
1020	x		x	x		x		x		x		x	x	
1033	x		x		x	x		x		x		x	x	
1038	x		x			x	x	x		x			x	
1046		x	x	x		x		x		x	x		x	
1050	x	x	x	x	x	x	x	x	x	x	x	x	x	x
1058	x	x	x		x	x		x	x	x		x	x	
1068	x	x		x		x		x		x			x	
1076	x	x	x		x		x	x		x		x		x

Figure 15: Verification example

To verify the model, we used non-random sampling in combination with the backtracking with the unique identifiers per delivery unit. This means that we first aggregate the data into overviews that give insights in what the result of the model is. We then search for interesting deviations from the current situation and check through all the spreadsheets how the model resulted in this deviation. For example, we see that after optimization, customer 1068 has a larger gap between the deliveries than in the current situation. We verified which delivery units were moved from and to which routes. Furthermore, we checked whether all restrictions regarding capacity and shelf life were met. Because we use a unique identifier for each delivery unit and we use this unique identifier in all spreadsheets and in the RVRP solver, we check exactly which delivery units are transferred from delivery moment. We know for each delivery unit exactly what restrictions apply regarding shelf life, and delivery window. Manual calculations as well as input/output cross checking has led to the belief that the model holds the restrictions that we inputted.

Validation

Because we only consider the eastern part of the Netherlands, the validation process cannot be done by exactly comparing the base case scenario within Paragon with reality. We do not have the same routes as HAVI drives in practice. We included a set of customers in our scope, instead of a set of routes. To validate the routing schedule that Paragon proposes we asked the transport planners to check whether the schedule is realistic. From practice HAVI already knows that the error in the times proposed is in the range of 20 minutes on average per route. The main concern the transport planners have on the proposed schedule is that the routes in the evening shift all depart quite early. It is Paragons nature to plan everything as early as possible. However, in practice HAVI wants the day routes to depart as early as possible and the evening routes as late as possible, to have a buffer between the two shifts. Furthermore, we found that Paragon can deliver one customer twice in a route, although one of the settings we putted into Paragon is that deliveries at the same customer should always be 18 hours apart. The situation was such that a delivery unit must be delivered before 9 am and another delivery unit must be delivered after 1 pm. In between the truck had to go to other customers. This happened only occasionally, and when seen, we tried to manually place the delivery unit in another route, by making as little changes as possible to the rest of the schedule. To compare the results of the various experiments, we compare with the base case scenario, which is the set of customers we selected, rerouted as if it is the total network.

4.9 Operational policies

The second model is for the operational policy. This policy is used to test which results are achieved with the new created routing schedule when volumes deviate from the forecasted volumes. We first give an overview of the operational policy with the requirements. Then we determine the minimum and maximum quantity that can be delivered. Based on the utilization of the routes we determine which part of volume should be added or removed from every delivery.

4.9.1 Operational policy overview

The tactical routing schedule is optimized using deterministic volumes. The volumes used in our model, are the same volumes that are being used by HAVI to create the tactical routing schedule that is driven from the 6th of July 2020. To test the performance of the tactical routing schedule with deviating volumes (i.e., performing a sensitivity analysis), we apply volumes of three different weeks. One week with a lower (week 22), one week equal to (week 27) and one week above (week 30) forecasted volume. Most important is that, although we know in advance which volume is input for the tests, we create the policy such that we only act upon data that was available at that time.

We test the operational policy on the current situation as well as on the best created new tactical routing schedule in which we didn't apply intervention (v). This operational policy must describe the following:

- How to determine the quantity to order for a delivery?
- How to re-optimize the routes of each day?

Using the following data:

- Actual sales per restaurant per day (to determine inventory levels).
- Forecasted volumes (to forecast utilization of routes of coming days).
- Hourly sales distribution (to determine which part of sales of a day took place before delivery).
- Capacities of storage locations as determined in Section 4.4 and 4.6.

Before we start with importing the volumes of the three chosen weeks, we decide which initial inventory levels we assume. Therefore, we add another week that we execute first, which results in an ending inventory, which is the starting inventory of the following week. We thus execute four weeks, and the ending inventory of a week determines the starting inventory of the following week.

For the start-up period, 7 days are used, followed by 21 days from the three weeks that are used for the operational results. Because we only act upon the data that was also available at that time, we do the calculations chronologically over time. For every day we execute the algorithm of Section 4.9.2.

4.9.2 Algorithm of the operational policy

In contrast to the rest of this research we assume for the operational policy, that we do not have to deliver full delivery units. For the first day we assume the starting inventory being equal to 50% of the capacity of the storage. For all other days, the starting inventory is equal to the ending inventory of the prior day. Therefore, the inventory for customer C and temperature zone T at the start of day d is as follows:

Equation 8

$$I_{C,T,d}^{start} = \begin{cases} 0,5 * Cap_{C,T} & for\ d = 1 \\ I_{C,T,d-1}^{end} & for\ d > 1 \end{cases}$$

The expected sales for customer C at day d is denoted by $ES_{C,d}$. If that customer gets a delivery at day d, the expected sales that takes place before delivery is denoted by $ES_{C,d}^{before}$. This data is retrieved from the data as presented in Table 4. To determine the inventory level just before delivery takes place, we use a constant that states how many sales corresponds to the consumption of one delivery unit. Assuming this constant to be equal for all temperature zones and is denoted by SU. For each customer the ratio of delivered units in each temperature zone is determined. We denote the percentage of delivered units for customer C that belongs to temperature zone T by $P_{C,T}$. From this, we determine for customer C that gets a delivery at day d what the expected consumption at day d is before delivery takes place from temperature zone T. Which is denoted as follows:

Equation 9

$$Cons_{C,d}^{before} = \frac{ES_{C,d}^{before}}{SU} * P_{C,T}$$

The inventory level just before delivery takes place at day d for customer C and temperature zone T is as follows:

Equation 10

$$I_{C,T,d}^{before} = I_{C,T,d}^{start} - \frac{ES_{C,d}^{before}}{SU} * P_{C,T}$$

Using the source data of Table 4, the expected sales between the delivery at day d for customer C and the next delivery and is denoted by $SN_{C,d}$. The consumption of goods belonging to temperature zone T for customer C between the delivery at day d and the next delivery is as follows:

Equation 11

$$CN_{C,T,d} = \frac{SN_{C,d}}{SU} * P_{C,T}$$

To determine the maximum quantity that can be delivered we must consider the capacity of the storage as well as the shelf life. Considering the capacity of the storage, the maximum quantity that can be delivered at day d to customer C for temperature zone T is as follows:

Equation 12

$$Q_{C,T,d}^{min} = CN_{C,T,d} - (I_{C,T,d}^{start} - Cons_{C,d}^{before})$$

To determine the maximum quantity that can be delivered we must consider the capacity of the storage as well as the shelf life. As we consider the capacity of the storage, the maximum quantity that can be delivered at day d to customer C for temperature zone T is as follows:

Equation 13

$$Q_{C,T,d}^{max,cap} = Cap_{C,T} - (I_{C,T,d}^{start} - Cons_{C,d}^{before})$$

For the shelf life criteria, we need to know the expected sales during the period of the shelf life agreement. As mentioned in Section 4.5, for the chilled goods, the agreement is four days, for the other temperature zones we set the shelf life at seven days. We denote the expected sales for temperature zone T for customer C during the shelf life period (including the whole day of delivery) by $SS_{C,T}$. Considering the shelf life, the maximum quantity that can be delivered at day d to customer C for temperature zone T is as follows:

Equation 14

$$Q_{C,T,d}^{max,shelf} = \frac{SS_{C,T}}{SU} * P_{C,T}$$

We conclude, the maximum quantity that can be delivered at day d to customer C for temperature zone T is as follows:

Equation 15

$$Q_{C,T,d}^{max} = min(Q_{C,T,d}^{max,cap}, Q_{C,T,d}^{max,shelf})$$

To determine the quantities that we are going to deliver, a target is set. This target is based on the quantity, which is used in making the tactical routing schedule. We adjust this quantity such that it fits between the minimum and maximum quantity determined in Equation 12 and Equation 15. The procedure is slightly different for the base case and the new tactical routing schedule.

Base case

The base case scenario is built in the same way as HAVI creates their tactical routing schedules in practice. This means that when a delivery quantity is determined, this is not specifically allocated to temperature zones. Therefore, we assume in the base case, that for all deliveries the delivered quantity per temperature zone is as follows:

Equation 16

$$Q_{C,T,d} = P_{C,T} * Q_{C,d}$$

In which $Q_{C,d}$ is the delivered quantity for customer C at day d and $Q_{C,T,d}$ is the delivered quantity for customer C for temperature zone T at day d.

New tactical routing schedule

In the new tactical routing schedule, we used a list of specific delivery units to deliver. For each delivery unit we know to which temperature zone it belongs. Therefore, we know exactly how much delivery units per temperature zone are included in a delivery and this is denoted by $Q_{C,T,d}$. We adjust the target quantity to fit between the minimum and maximum quantity as follows:

Equation 17

$$T_{C,T,d} = min(max(0\,, Q_{C,T,d}\,, Q_{C,T,d}^{min})\,, Q_{C,T,d}^{max})$$

For all other customers than McDonald's we use the target quantity $T_{C,T,d}$ being equal to the average of the past 3 weeks that is delivered to that customer on the specific day of the week, increased by 10%. Because we know which customers are delivered together in a route we sum all adjusted targets $T_{C,T,d}$ that are delivered in route r to obtain the target level of the route denoted by TR_r.

Optimizing utilization

At this point the goal is to utilize the existing routes as much as possible. All routes have a maximum capacity of sixty delivery units. Using the set targets, route r is filled with TR_r delivery units. If $TR_r < 60$ there is overcapacity and if $TR_r > 60$ there is capacity shortage.

Overcapacity

If $TR_r < 60$, there is space left in the truck, which we fill by pulling some delivery units from following deliveries. To determine which delivery units to add to route r, we determine for all deliveries within that route (on day d) if there is overcapacity for customer C and temperature zone T and this is denoted as follows:

Equation 18

$$OC_{C,T,d} = Q_{C,T,d}^{max} - T_{C,Td}$$

We take the sum of all temperature zones and all customers that are being delivered in route r to obtain the total quantity of delivery units that can be added to route r and denoted by OC_r. Important to notice that OC_r has nothing to do with the capacity that is left in a truck, but it is the capacity that is left within the storages of the customers. The fraction of the overcapacity for the customers and temperature zones used in route r is denoted by f_r and determined as follows:

$$f_r = min(\frac{(60 - TR_r)}{OC_r}\,, 1)$$

The fraction used cannot be larger than one because then we would have delivered more than the maximum quantity as determined in Equation 15. We determine the number of delivery units that we deliver to customer C for temperature zone T at day d as follows:

Equation 19

$$D_{C,T,d} = T_{C,T,d} + f_r * OC_{C,T,d}$$

Capacity shortage

If $TR_r > 60$, we have a shortage of space in the truck, which needs to be freed if possible. The calculation has an analogy with the scenario of overcapacity. Some variables are the same, but this is possible because a route can never have overcapacity and capacity shortage at the same time. To determine which delivery units from route r we must delay to the following delivery, we determine for all deliveries within that route (on day d) if there are delivery units for customer C and temperature zone T that can be delayed and denoted:

Equation 20

$$DD_{C,T,d} = T_{C,T,d} - Q_{C,T,d}^{min}$$

We take the sum of all temperature zones and all customers that are being delivered in route r to obtain the maximum quantity of delivery units that can be delayed from route r and denoted by DD_r. The fraction of the delivery units that can be delayed for the customers and temperature zones in route r that we use is denoted by f_r and determined as follows:

Equation 21

$$f_r = min(\frac{(TR_r - 60)}{DD_r}, 1)$$

The fraction used cannot be larger than one because then we would have delivered less than the minimum quantity as determined in Equation 12. We determine the number of delivery units that we deliver to customer C for temperature zone T at day d as follows:

Equation 22

$$D_{C,T,d} = T_{C,T,d} - f_r * DD_{C,T,d}$$

So far, we only used expected sales to determine the delivery quantities. Before the process repeats, we determine the end inventory of the day based on the actual sales. The actual sales of customer C at day d is denoted by $AS_{C,d}$. The inventory at customer C in temperature zone T at the end of day d is denoted as follows:

Equation 23

$$I_{C,T,d}^{end} = I_{C,T,d}^{start} + D_{C,T,d}$$

It is possible that $I_{C,T,d}^{end}$ becomes negative, then there is a stockout. In this case an extra delivery needs to be created that must be delivered at day d to customer C with $-I_{C,T,d}^{end}$ delivery units in temperature zone T. Including this extra delivery $I_{C,T,d}^{end}$ is set equal to zero. In practice, this volume is not accurate because we must consider $Cons_{C,d}^{before}$, which increases the size of the extra delivery. We also must consider the fact that a buffer storage is in place, which reduces the size of the extra delivery.

For every day for which we execute the operational policy, we repeat everything in Section 4.9.2. The policy described, determines for every delivery how much delivery units to deliver.

We advise HAVI to develop/adjust the system to automatically create the orders with the above determined quantities. Further research is needed to translate the model to the level of products instead on delivery units. For the time being, the restaurant planners adjust the orders, so they comply with the determined volumes. The orders made by the system of HAVI are based on JIT delivery and match the minimum quantities as computed in Equation 12.

4.9.3 Optimizing the operational routing, using the RVRP solver

In Section 4.9.2, we determined the delivery quantities for the operational deliveries. The order quantities are always determined the day before they are delivered. After the confirmation of the orders, the transport planners optimize the routes. They aim to (i) minimize the number of adjustments and (ii) avoid inefficiencies. At this moment, the customers that get a delivery are fixed and the delivery quantities are being fixed. Therefore, we solve a VRP problem.

In practice, the transport planners solve this problem manually, by transferring deliveries between routes within a single day, for the deliveries of tomorrow. In this research, we use Paragon, our RVRP solver, to have the most consistent approach. We cannot minimize the number of adjustments made, but we can limit the complexity of adjustments that Paragon can make. The more complex the adjustments that are allowed, the better the result, at the expense of calculation time. With our settings the calculation time is on average shorter than five seconds. Furthermore, we limit all deliveries to be delivered within three hours of the delivery time in the tactical routing schedule. The time windows from the master data still apply.

We advise HAVI to do further research to adjust the routes a few days in advance. With help of the operational policy and forecasted volumes, we can predict which routes will be overloaded even after the operational policy will be executed. By acting a few days in advance, it is managed where extra volume should be delivered. If for example multiple routes on Friday are overloaded and we anticipate early enough, we can add a route or extra deliveries on Thursday. This way more volume could be delivered on Thursday instead of Friday.

4.10 Conclusion

In this chapter we answered research questions 3, 4, 5, 6b and 7.

3. *How can the requirements be translated into restrictions/input that can be used in routing optimization software?*

Some of the requirements are met by using some settings within Paragon. Paragon uses time windows, optimizes based on 'clustering based on duty time' and adds extra duty time to a route to have a buffer time between two shifts. Setting the priorities can be used in the same way to minimize workload on the days to which salary supplements apply. Shelf life restrictions and storage restrictions are included by Equation 2 and Equation 3 respectively. The reduction in operational costs and operational adjustments needed will be tested in Chapter 5.

4. *How can workload balancing be applied in the tactical routing process of HAVI?*
    a. *What is the objective of HAVI that should be used in this research?*

Workload balancing can be achieved when HAVI takes the lead in determining the delivery patterns instead of following the wishes of the customers. Balancing the workload is made easier by planning every delivery unit instead of planning deliveries. The allocation procedure can also be used to balance the workload by giving the highest priorities to the shifts with the lowest workload. Intervention (ii), (iii) and (iv) promote workload balancing. Dummy customers are used to prevent routes to transfer to shifts with lower priority.

b. *What input is required to set up the model that will be implemented into the processes of HAVI?*

The most important input for our model is the storage capacity per temperature zone per restaurant. These capacities are deduced based on two criteria, which determine a lower bound. To calculate this lower bound we used three sets of data:

- Historical deliveries data.
- Forecast for the near future.
- Sales distributions.

After determining the lower bound based on these criteria, we used the fact that the current routing schedule consists of feasible delivery patterns to relax the lower bounds further, such that for every delivery moment there is at least one possible delivery moment.

5. *How do we build the model to optimize the transport and inventory integrated?*
    a. *How does our model differentiate relative to the known models from literature?*

Some of the main differences between our model and the traditional approaches from literature are mentioned earlier and are stated in Table 2. Our model includes the following characteristics, which are not frequently covered in the known models from the literature:

- Delivery quantity is dependent on the exact timing of the delivery, not only the delivery day.
- Multiple groups of products are considered in storage as well as shelf life.
- The delivery patterns are optimized simultaneous with the routing process.
- Workload balancing is considered.
- Salary supplements are used, and a tradeoff can be made with balancing workload.

b. *Which steps are used in the model?*

This research is unique because we determine the delivery patterns and delivery quantities in a single phase (see Section 4.2) in contrast to almost all models in literature who determine the two aspects in different optimization phases. We first determine for each delivery unit the delivery window (see Section 4.5), based on the expected hourly sales given a historical distribution, considering shelf life agreements and the capacity of storages. We built an allocation procedure (see Section 4.7.2), which we use to steer on which shifts to utilize most. Furthermore, we built an algorithm consisting of nine steps (see Section 4.7.2). Step 1 to 7 is heuristic A, which aims at removing unnecessary workload in shifts with low priorities. Step 8 and 9 is heuristic B, which aims at utilizing the routes in the shifts with high priorities and reducing workload in shifts with low priorities.

6. How can the tactical model be tested using experiments?
    b. How to compare the results of the experiments?

The main goals following from the main research question are (i) reducing costs and (ii) improving the balance of the workload. Therefore, we compare the experiments mainly on these aspects. How the costs and the balancing of the workload are measured is explained in 4.7.1. In Chapter 5, the main results are discussed. Explanatory results are stated in Appendix X.

7. *What operational policy should be followed after implementing the new tactical routing schedule?*
    a. *What operational policy should the restaurant planners follow?*

In Section 4.9.2 we first determine the minimum and maximum quantity that can be delivered. We use the targets, which are based on the tactical routing schedule to determine whether a route is planned above or under capacity. We then adjust the targets such that we more utilize the capacity of the trucks or postpone delivery units where possible if a truck is planned above capacity.

The system that HAVI uses calculates the minimum quantities that should be ordered. The restaurant planners should adjust the orders by pulling volume from subsequent deliveries, such that the order quantity is increased to the volumes determined in Equation 22.

    b. *What operational policy should the transport planners follow?*

The transport planners solve VRP problems after orders are being confirmed. We advise to minimize the number of deliveries to switch from shift. Further research must be done to anticipate on overloading routes, so HAVI can manage where extra volume should be delivered.

5. Experiments

In this chapter we answer research questions 6a and 8. These research questions are as follows:

Research question 6a: '*Which scenarios must be tested?*'

Research question 8: '*What performance can be achieved by applying the new tactical model and operational policies into practice at HAVI?*'

In Chapter 4 we built the model that we use to improve the current tactical routing schedule. In Section 5.1 we determine which experiments to execute. The experiments differ on multiple factors:

- Primary focus on either cost reduction or balancing workload.
- Choice of heuristics applied in intervention (i).
- Either applying interventions (ii) to (v) or not.

The experiments are compared with each other in Section 5.2 based on the following criteria:

- Routes per shift.
- Delivery units per shift and warehouse related costs.
- Utilization per shift.
- Total duty time per shift and corresponding costs.
- Total traveled distance per shift and corresponding costs.
- Impact on fleet size and corresponding costs.
- Average delivery size per shift.
- Centre of Gravity.

The criteria above are explanatory results, which influence the main results which are:

- Total costs (costs for trucks, fuel and labor).
- Contribution to balancing the workload in the total network.

The main results are mentioned in this chapter. The rest of results are used to support the findings and are mentioned in Appendix X. From all the experiments, we choose the best tactical routing schedule which we use to perform a sensitivity analysis in Section 5.3 using operational data, with historical volumes.

5.1 Experimental design

Before elaborating on the experiments, we discuss the base case, which is used as a reference.

5.1.1 Base case scenario

Unfortunately, we cannot take the current situation as the base case scenario. At the expense of computation and research time, we choose to only include the eastern part of the Netherlands in this research, which contains about 50 restaurants of McDonald's and 29 locations of other customers. We take all routes from the current routing schedule that contains any of these customers. We inserted the routes with only the in-scope customers into Paragon. We inserted deliveries with fixed volumes, which are also used to build the current tactical routing schedule. Because we only imported the customer in scope, some of the routes had a low utilization. We re-planned the routing schedule as if it were the complete network we were looking at, and as if it were going to be the next routing schedule that would be implemented by HAVI. We now have a tactical routing schedule, with the same delivery patterns as in the current situation, but with other routes.

5.1.2 Experiments

In Table 9 we show the experiments we execute. The third column is used to state the choice made in the first intervention. The fourth column is used to state all other interventions. Intervention (iii) and (iv) can be executed multiple times in row. We choose to execute an intervention at most two times. The experiments are executed in three stages, in Table 9 separated by black lines. The experiments in the first stage are fully determined in advance. The experiments in the second stage are only executed for the best option of intervention (i) per primary focus. The third stage consists of an experiment in which we will add intervention (v) to all experiments that are on the efficient frontier (measuring costs vs. balancing workload). The results that support these decisions are discussed in the next section. The calculation times of the experiments can differ.

Table 9: List of experiments

Experiment	Primary focus	Intervention (i)	Interventions (ii) to (v)
I	Cost reduction	A	
II	Cost reduction	B	
III	Cost reduction	A & B	
IV	Balance workload	A	
V	Balance workload	B	
VI	Balance workload	A & B	
VII	Cost reduction	A & B	(ii)
VIII	Cost reduction	A & B	(iii)
IX	Cost reduction	A & B	2x (iii)
X	Cost reduction	A & B	(iv)
XI	Cost reduction	A & B	2x (iv)
XII	Balance workload	B	(ii)
XIII	Balance workload	B	(iii)
XIV	Balance workload	B	2x (iii)
XV	Balance workload	B	(iv)
XVI	Balance workload	B	2x (iv)
XVII	Cost reduction	A & B	2x (iv) & (v)

5.2 Experimental results tactical

In this section we discuss the results of the tactical routing schedule. We first discuss the performance of the experiments in terms of running time. We then discuss the performance of the experiments in terms of output.

The experiments vary large with respect to running time (i.e., CPU time needed). Executing the experiments without interventions we needed between five and thirty minutes CPU time. By applying the interventions, the CPU time is most influenced by the number of times we optimized the routing schedule using the RVRP solver. For example, in the experiment where we reduce the fleet size with two vehicles, we have to remove deliveries until it is possible to save a route in each of the required shifts. In experiment XVI we had to remove 27 deliveries before we were able to save a route in each of the required shifts. That also means 27 times running the RVRP solver, which has cost a whole day, of which estimated 4 hours CPU time.

The main results of the experiments are the total costs of the tactical routing schedules and how they contribute to balancing the workload of the total network, expressed in a penalty (the lower, the better). In Section 4.7.1 we discussed how these aspects are computed. We calculate for all the experiments the costs and the workload penalty WP_e. The results are stated in Table 10 and visualized in Figure 16.

From the first three experiments it resulted that experiment III was the best performing on costs as well as workload balancing. Therefore, the rest of the experiments with a primary focus on cost reduction are executed with heuristic A & B. From the fourth to sixth experiments, experiment V performs better on costs as well as balancing workload compared to experiment IV. Experiment V performed much better than experiment VI on costs, but slightly less on workload balancing. Therefore, the rest of the experiments with a primary focus on workload balancing are executed with heuristic B. From all experiments from the first two stages, experiment XI performed best on costs as well as workload balancing. Therefore, experiment XVII is the same as experiment XI with additional intervention (v).

We see in Figure 16 that the efficient frontier, the green circle, consists of only one experiment. Looking at the trend line, we see that in general, costs and workload balancing is not a tradeoff but is correlated positively to each other. The correlation factor $r(Costs, WP)$ is computed as follows:

Equation 24

$$r(Costs, WP) = \frac{\sum_e (Costs_e - \mu_C) * (WP_e - \mu_{WP})}{\sqrt{\sum_e (Costs_e - \mu_C)^2 * \sum_e (WP_e - \mu_{WP})^2)}} = 0.760$$

in which $Costs_e$ are the costs related to experiment e, μ_C are the average costs over all experiments and μ_{WP} is the average workload penalty over all experiments. To test for significance of the correlation we compute the t value as follows, using an α of 0.05:

H_0: The null hypothesis states that the correlation factor r is equal to zero.

H_1: The alternate hypothesis states that the average costs when focusing on costs first, is lower than the average costs when focusing on balancing workload first.

Equation 25

$$t = \frac{r * \sqrt{n-2}}{\sqrt{1-r^2}} = 4.68$$

The critical $t - value$ with 16 (i.e., $n - 2$, in which n is the number of experiments) degrees of freedom is 2.12. Because the $t - value$ we found is higher than the critical t-value, we reject the null hypothesis and conclude that the correlation between costs and workload balancing is significant.

Table 10: Experiments main results

Experiment	Total €	WP^e
Base	€ 38.615	234,0
I	€ 39.585	259,8
II	€ 40.869	248,2
III	**€ 37.713**	**232,9**
IV	€ 38.742	252,7
V	**€ 38.008**	**241,3**
VI	€ 39.686	240,9
VII	€ 37.359	238,1
VIII	€ 36.758	239,8
IX	€ 38.146	234,9
X	€ 35.686	243,2
XI	**€ 33.395**	**215,7**
XII	€ 38.513	225,0
XIII	€ 37.574	241,1
XIV	€ 37.866	243,3
XV	€ 35.744	236,7
XVI	€ 33.608	218,6
XVII	**€ 32.641**	**214,9**

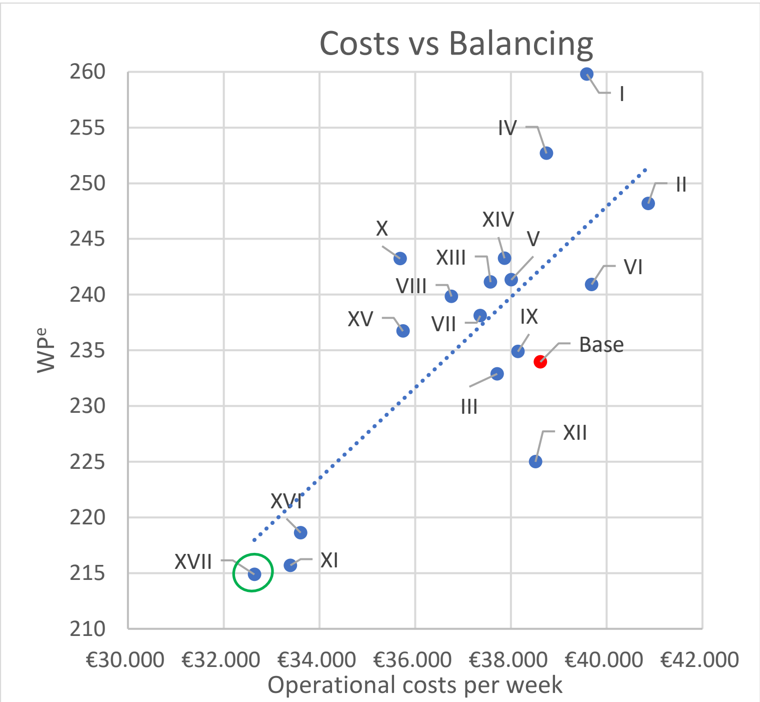

Figure 16: Costs vs. Workload balancing

The best tactical routing schedules are in the bottom-left part of Figure 16. The three experiments with the lowest costs and penalties all have in common that they have applied intervention (iv) twice and thus the fleet size in the total network is reduced with two trucks. In Section 5.2.2 we test the significance of the fleet size on the results.

The best routing schedule without intervention (v) is experiment XI, in which €5007 is saved compared to the base case. With intervention (v) an additional saving of €967 can be achieved. In experiment XVII are 10 deliveries that have a consecutive delivery within 18 hours.

Interpreting the results, we see in Table 39 in Appendix X, that the costs directly related to the fleet size vary from - €2000 to +€4000 , which means a spread of €6000. The warehouse related costs have a spread of €590 as is shown in Table 34 in Appendix X. The labor costs related to transport have a spread of €1351 as is shown in Table 36 in Appendix X. The costs related to the distance traveled have a spread of €1364 as is shown in Table 37 in Appendix X. The spread of costs related to the fleet size has a larger spread, than all other spreads summed together. Therefore, the fleet size has the largest impact on the costs. Another explanatory result is the utilization of the trucks, which is, according to Song and Savelsbergh (2007), often a good indicator for the transportation costs, given the volume that is delivered. In Table 35 in Appendix X, we see that in experiment XI, the utilization has increased to 89% compared to 86% in the base case.

5.2.1 Paired T-tests

To test whether the primary focus of the algorithm (i.e., either we focus on cost reduction first, or we focus on balancing workload first) has the intended outcome is verified with paired T-tests. In Table 11 we state which experiments are paired for the T-test. We use an α of 0.05. We assumed equal variances in the sets.

Costs

H_0: The null hypothesis states that the average costs when focusing on costs first, is equal to the average costs when focusing on balancing workload first.

Table 11: Pairs from T-test

Experiment Costs	Experiment Balancing workload
I	IV
II	V
III	VI
VII	XII
VIII	XIII
IX	XIV
X	XV
XI	XVI

H_1: The alternate hypothesis states that the average costs when focusing on costs first, is lower than the average costs when focusing on balancing workload first.

Applying the T-test on the costs, we find a P-value of 0.49, meaning that we cannot reject the null hypothesis. Therefore, there is no significant difference between the average costs when focusing on costs first and the average costs when focusing on balancing workload first.

Balancing workload

H_0: The null hypothesis states that the average WP_e when focusing on balancing workload first, is equal to the average WP_e when focusing on costs first.

H_1: The alternate hypothesis states that the average WP_e when focusing on balancing workload first, is lower than the average WP_e when focusing on costs first.

Applying the T-test on the costs, we find a P-value of 0.39, meaning that we cannot reject the null hypothesis. Therefore, there is no significant difference between the WP_e when focusing on costs first and the WP_e when focusing on balancing workload first.

Concluding, we show that the various settings of priorities as stated in Appendix VIII do not have a significant impact on the relation between costs and workload balancing.

5.2.2 One-way ANOVA

To verify whether the fleet size has indeed a significant impact on our results we perform One-way ANOVA test. We use an α of 0.05.

Costs

H_0: The null hypothesis states that the average costs do not depend on the fleet size.

H_1: The alternate hypothesis states that the average costs depend on the fleet size.

With a one-way ANOVA we do not find which of the sample means differs significantly from others. We find the following results:

Table 12: One-way ANOVA, fleet size impact on costs

Impact fleet size	*Count*	*Sum*	*Average*	*Variance*
-2	3	99643	33214	257921
-1	2	71430	35715	1711
0	11	418879	38080	589102
+1	2	80554	40277	699449

Performing a variance analysis, we find an F-value of 52.2 with a critical F-value 3.34. Therefore, we reject the null hypothesis. We show that there is a significant relation between fleet size and the total costs. Table 12 shows that the costs rise when the fleet size increases. In our research we cannot decrease the fleet size with more than two trucks. In the total network the largest shift is Friday daytime. During that shift, in the eastern part of the Netherlands three trucks are used in the current situation. Some of the delivery units must be delivered on Friday daytime, due to their small delivery windows. Therefore, at least one route should be remained in the east.

In Table 13 we see that even when we exclude the costs directly related to the fleet size, the costs still increase with the fleet size. However, we find an F-value of 1.39 with a critical F value of 3.34. Therefore, we cannot reject the null hypothesis. The relation of fleet size to the operational costs, where costs directly related to the fleet size are excluded, is not shown to be significant.

Table 13: One0way ANOVA, fleet size impact on operational costs

Impact fleet size	*Count*	*Sum*	*Average*	*Variance*
-2	3	111643	37214	257921
-1	2	75430	37715	1711
0	11	418879	38080	589102
+1	2	76554	38277	699449

Workload balancing

H_0: The null hypothesis states that the average WP_e does not depend on the fleet size.

H_1: The alternate hypothesis states that the average WP_e depends on the fleet size. We find the following results:

Table 14: One-way ANOVA, fleet size impact on balancing workload

Impact fleet size	*Count*	*Sum*	*Average*	*Variance*
-2	3	649	216	3,88
-1	2	480	240	21,2
0	11	2643	240	90,8
+1	2	489	245	26,4

Performing a variance analysis, we find an F-value of 7.36 with a critical F-value 3.34. Therefore, we reject the null hypothesis. We show that there is a significant relation between fleet size and workload balancing. Table 14 shows that the WP_e rises when the fleet size increases.

5.2.3 Linear regression

We use linear regression models to test the significance of the interventions on the results. We first test for the impact on costs. We use an α of 0.05. The result of the linear regression is stated in Table 15. We see that only intervention (iv) has a significant impact on reducing the costs (i.e., intervention (iv) is the only intervention having a P-value below 0.05).

Table 15: Significance of impact interventions on costs

intervention	*Coefficient*	*Standard error*	*t Stat*	*P-value*	*Lower 95%*	*Upper 95%*
ii	-932	780	-1,20	0,25	-2631	766
iii	-685	379	-1,81	0,10	-1511	140
iv	-2778	379	-7,33	0,00001	-3603	-1952
v	-672	1149	-0,58	0,57	-3176	1832

We also test the impact of the interventions on the workload balancing measured by WP_e. The result of the linear regression is stated in Table 16. We see that intervention (ii) and (iv) have a significant impact on reducing the penalty for workload balancing.

Table 16: Significance of impact interventions on WP

intervention	*Coefficient*	*Standard error*	*t Stat*	*P-value*	*Lower 95%*	*Upper 95%*
ii	-15,2	6,24	-2,43	0,032	-28,8	-1,57
iii	-4,30	3,03	-1,42	0,18	-10,9	2,31
iv	-13,2	3,03	-4,34	0,001	-19,8	-6,55
v	-5,50	9,20	-0,60	0,56	-25,5	14,5

5.3 Operational results

The operational policy is applied to the tactical routing schedule that we created with experiment XI, to test the following:

- Number of shifted routes.
- Number of changed delivery times > 30 minutes.
- Number of extra deliveries needed.
- Impact on costs.
- Sensitivity analysis with deviating volumes.

The results from the criteria above are discussed in this section. The explanatory results are stated in Table 40 to Table 46 in Appendix X. The tactical routing schedule is based on a forecast with a total of 2567 delivery units per week. The operational policy is applied on three different weeks. The volumes measured in delivered delivery units per week are stated in Table 17.

Table 17: Weeks used for testing operational policy

Week	Number of delivery units	Deviation from forecast
Forecast (see Section 4.3)	2567	
22 (25/05/2020 – 31/05/2020)	2357	- 8,2%
27 (29/06/2020 – 05/07/2020)	2596	+1,0%
30 (20/07/2020 – 26/07/2020)	2799	+8,0%

Shifted routes

Counting the number of shifted routes is important for the workforce planning. When the number of routes in a shift deviates from the planning, drivers are cancelled or arranged lastminute. The number of changes is stated in Table 18.

Table 18: Number of shifted routes

Week	Current situation	New routing schedule
22	1	8
27	0	9
30	1	13

According to Table 18, the new routing schedule is more unstable when looking at the number of changes that are made in the workforce scheduling. Also, when volume is closest to forecasted volume, nine changes must be made. We advise HAVI, when optimizing the operational transport planning, to minimize the number of deliveries to switch from one shift to another. In practice, the transport planners rarely switch deliveries from shifts. While using the RVRP solver, we did not restrict deliveries to be switched from shift. Therefore, in the new routing schedule more deliveries are delivered in another shift than they were planned in the tactical routing schedule. The base case is made as much as possible in the same way as the tactical routing schedules made in practice by HAVI. Evening routes depart such that the last delivery takes place around 11 pm. In the new situation the RVRP solver had the freedom to choose the departure times. Paragon always schedules everything as early as possible. The average departure time of evening routes in the current situation is 4:40 pm. The average departure time of evening routes in the new routing schedule is 2:51 pm. This contributes to the ability to switch deliveries from shift.

Changed delivery times

When volume deviates from forecast, is it necessary to change some routes to either make routes feasible or gain efficiencies. In Table 19, we count deliveries that are changed more than 30 minutes. The current situation is not measuring how many changed delivery times actually have taken place, but how many changed delivery times take place if the operational policy we created is applied to the base case routing schedule.

Table 19: Number of changed delivery times (greater than 30 minutes)

Week	Current situation	New routing schedule
22	54	58
27	65	67
30	49	113

From this result we conclude that the number of changed delivery times is not much impacted, unless volumes are higher than forecasted. Remarkable is that the RVRP solver shifted some routes from Tuesday evening to Tuesday daytime in the high-volume week, for the new schedule, as is shown in Table 40 in Appendix X. The underlying reason for this shift can be concluded from Table 42 in Appendix X where we see that the utilization of Tuesday daytime is 99% and Tuesday evening 68% with forecasted volumes. When volumes increase, the routes on Tuesday daytime overflow, but Tuesday evening has still spare capacity. Therefore, some of the deliveries originally delivered in the Tuesday daytime routes, should be combined with some deliveries from the Tuesday evening routes. Apparently, in this specific example, it was most efficient to plan those routes in the day shift. Intuitively, we argue that the utilization of the truck has an influence on the number of changes made in number of routes per shift and delivery times. In the new routing schedule, the utilization is increased to 89% compared to 86% in the base case as is shown in Table 35 in Appendix X.

Stockouts

When a stockout occurs, an extra delivery is needed to replenish the storage with enough goods until the next regular delivery. On any given moment, it can be discovered that the storage at a specific customer is not enough to meet demand until the next delivery. We assume that we know in time enough that a stockout occurs, so we are able to schedule the extra deliveries in regular routes. If in practice the stockout is discovered too late, a carrier is needed who drives dedicated to this customer. The number of extra deliveries needed are stated in Table 20. These extra deliveries are also given as input for the RVRP solver, and thus the extra deliveries are already included in al results.

Table 20: Number of extra deliveries needed

Week	Current situation	New routing schedule
22	0	0
27	0	3
30	3	15

With higher volumes, more stockouts occur in the new tactical routing schedule. Intuitively this is because in the new tactical routing schedule we determined the boundaries using a 110% scenario. However, when the market has a plus of 8%, some specific customers will have sales plusses of more than 10%. The customers, for whom the delivery pattern is chosen close to the boundaries, will face a stockout in this situation. It is an exceptional situation that in week 30 sales were 8% above forecast.

Due to Covid-19, volumes are harder to predict. Further research should be done to test the impact of determining the delivery windows based on other volumes instead of 110%. The correlation can be found between total costs and stockout occurrences.

Impact on costs

The costs associated with the experiment of a tactical routing schedule is an indicator for the operational costs. In practice volumes always deviate from forecast and thus costs will be different. In Table 21 we mention the costs after applying the operational policy.

Table 21: Costs applying the operational policy

Week	Volume deviation	Current situation	New routing schedule
Forecast		€38.615	€ 33.608
22	-8,2%	€ 35.747	€ 30.341
27	+1,0%	€ 37.677	€ 34.570
30	+8,0%	€ 40.079	€ 39.374

We see that the new routing schedule is more efficient in all scenarios. The lower the volume, the greater the difference. In Section 5.2 we found that the fleet size has the largest impact on costs. In the new tactical routing schedule, we forced the RVRP solver to decrease the fleet size with two trucks. In week 22, the same fleet size could be remained. In week 27 and 30, volumes increase and respectively one and two extra trucks are needed. In week 30 no more trucks were saved compared with the current situation. Therefore, costs of the new routing schedule are more sensitive to volume deviations compared to the current situation. Summarizing from earlier results, we conclude that this sensitivity is caused by the higher utilization in the new routing schedule together with the extra deliveries that have to delivered due to stockouts. Costs for delivering extra deliveries for extra deliveries can be more expensive than the average deliveries, because the extra deliveries are not considered in the tactical routing schedule. Therefore, the extra delivery can be located geographically inefficient.

5.4 Conclusion

In this chapter we answered research questions 6a and 8.

6. *How can the tactical model be tested using experiments?*
    a. *Which scenarios must be tested?*

We executed 17 experiments in which we varied the primary focus of the allocation procedure and the interventions applied. The experiments executed are listed in Table 9. The experiments are divided in three stages. The first six experiments are used to find the best combinations of primary focus and the choice of intervention (i). Experiment VII to XVI build upon the best experiments from I to VI, which are used to test the impacts of interventions (ii), (iii) and (iv). The last experiment is used to improve the best experiment so far by applying intervention (v).

8. *What performance can be achieved by applying the new tactical model and operational policies into practice at HAVI?*

The main performances measured, are the costs and the workload balancing. We find that these two measures are significantly positively correlated. Two out of the five interventions have a significant impact on the results. The intervention that added deliveries to dummy customers to prevent the balance of number of routes per shift to get worse, contributes significantly to improve the balance of the workload. Forcing the RVRP solver to reduce the fleet size contributes significantly to reducing costs as well as improving the balance of the workload. In the best tactical routing schedule, we reduced the fleet size with two vehicles, which is the maximum that can be achieved, because only the eastern part of the Netherlands is in scope.

The best tactical routing schedule in which consecutive deliveries are 18 hours apart, saves €5007 per week compared to the current situation. Allowing consecutive deliveries to take place within 18 hours, saves another €967 per week. The contribution of workload balancing for the total network is expressed in a penalty. How this penalty is computed is explained in Section 4.7.1. This penalty is in the best tactical routing schedule reduced from 234 to 218.6, and to 214.9 with consecutive deliveries within 18 hours allowed.

We find that there is no significant difference in the results obtained with different primary focuses. Either focusing on costs first or focusing on workload balancing first can both improve the costs and the workload balancing.

We performed a sensitivity analysis of the best-chosen tactical routing schedule. We choose the tactical routing schedule from experiment XI, in which we forced the RVRP solver to reduce the fleet size with two vehicles. This routing schedule has the best results without allowing consecutive deliveries within 18 hours apart. We applied the operational policy with deviating volumes from three different weeks from historical data. We find that the new routing schedule is more sensitive to volume deviations than the current routing schedule. The savings achieved, are shown in Table 21.

In the new routing schedule, more changes are made in delivery times (238 instead of 168) and number of routes per shift (30 instead of 2) in a period of three weeks. We advise HAVI to research whether these numbers will decrease when the transport planners minimize the number of deliveries that are switched to another shift. The number of stockouts are increased from 3 to 18 in a period of three weeks. We advise HAVI to do further research whether the number of stockouts can be reduced by determining the delivery windows based on higher volumes.

6. Conclusions and recommendations

In Section 6.1 we present the conclusions from our research and our contributions to theory. In Section 6.2 we mention our recommendations for HAVI and provide options for further research. Most importantly, in this chapter we answer the main research question of this research, which is as follows:

How can HAVI save on operational costs by balancing the workload within a week, without violating customers' restrictions including storage capacities at the customers' locations?

6.1 Conclusions

Background and literature

One of the main underlying problems is that transportation and inventory management are separated within HAVI. First, we performed a literature study. In this study, where we studied VRP and IRP models. On an operational level, HAVI deals with an IRP model with fixed routes. The order quantities must be determined. The transport planners then solve a VRP model every day. On a tactical level HAVI also deals with an IRP model. The delivery patterns and quantities are both decision variables. We were not able to identify the literature that jointly optimizes the routing problem and the inventory control. In one phase, the delivery patterns are determined. In the other phase the delivery quantities are determined. This is the same as HAVI operates in the current situation. However, in this research, we strongly believe in an approach in which these two phases are combined in a single phase.

Contributions to theory

With our model and research, we have built upon existing literature. We found some gaps in literature for which we provided solutions. We were not able to identify the literature that jointly optimizes the routing problem and the inventory control. The method we created is build such that the RVRP solver simultaneously optimizes the delivery patterns and delivery quantities. Restrictions regarding the capacities of the customers' storage locations are included. Therefore, our model integrated the routing problem with the inventory control. With spreadsheet models we have constructed the input that the RVRP solver requires. This input consists of all delivery units with a corresponding delivery window. Because this delivery window considers, the capacities, we designed the model such that an extended IRP can be solved by an RVRP solver. Because of these delivery windows, which we have determined for each delivery unit, we achieved to let the delivery quantities be dependent on the delivery times (hour). Furthermore, we have included shelf life restrictions, which is one of the main obstacles to apply basic IRP models in the food sector. The fact that the goods consumed are divided in multiple temperature is included. Shelf life is different for different temperature zones, which largely influences the delivery windows of the different goods. Furthermore, the goods are stored in different storages, which all have their own capacity. We have created an algorithm that is used to be able to check how a routing schedule would have performed when actual volumes deviate from forecasted volumes.

Requirements

The main research question states that the objective is to save on operational costs by balancing the workload. Therefore, the model built, measures costs as well as balancing the workload as objective. The main research question also states that the restrictions regarding the customers' storage locations should be respected. Our model includes the storage capacities as restrictions.

To determine the delivery patterns and delivery quantities in a single phase, we build a model that determines a delivery window for every single delivery unit. The delivery window considers: the storage capacities, shelf life agreements, consumption periods and hourly sales distributions.

Storage capacities

The storage capacities are not determined earlier by HAVI. Earlier attempts of HAVI to include the inventories capacity in their models resulted in the conclusion that the surface of storage locations is a poor indicator for their capacity to store goods. We use two criteria to deduce the capacity of the storages locations by determine lower bounds:

1. The capacity is at least as big as the largest delivery that is received in the past.
2. The capacity is at least as big as the largest amount that is expected to be consumed between two deliveries within the planning horizon.

The second criterium builds on the assumption that the restaurant planners are aware of the expected volumes and that their conclusion is that the delivery pattern in the current situation is feasible. Furthermore, we used the fact that the current routing schedule consists of feasible delivery patterns to relax the lower bounds such that every delivery unit has at least one possible delivery moment in the current situation.

Delivery windows

For every delivery unit, we determine when it will be consumed. Therefore, we use the hourly sales distribution per restaurant. This distribution is based on historical data.

To determine the earliest possible delivery moment regarding the capacity restriction, we use the consumption periods in combination with the capacities of the storage locations. Besides this capacity restriction, there is also a shelf life agreement, which limits the possible delivery moment given the period of consumption. These two limitations together determine the earliest possible delivery moment. The latest possible delivery moment dependents on the start of the consumption period. We then have determined the complete delivery window.

Improvement heuristic

The model built to optimize the tactical routing schedule is an improvement heuristic that takes the current routing schedule as initial solution. The model is split in heuristic A and heuristic B.

Allocation procedure

In both heuristic A and B, we use an allocation procedure to determine for each delivery units in which route it will be delivered. The allocation procedure is based on setting priorities to the shifts. The shifts with the highest priority, are the shifts we want to utilize the most. We experiment with different settings of the priorities. Either we focus on cost reduction first, or we focus on balancing workload first. We found that there is no significant difference in the results for the different approaches.

The goals of heuristic A are:

(i) Avoid unnecessary deliveries in the most costly or largest shifts (i.e., shifts with the lowest priorities).
(ii) Add deliveries to the shifts that are less costly or which workload is lower (shifts with highest priorities).

The purpose of heuristic B is to move the workload as much as possible to shifts with the highest priorities. In the allocation method the delivery units are assigned as much as possible to shifts with the highest priorities. The goals of heuristic B are:

(i) Use the opportunities to merge routes in the shifts with the lowest priorities because they have a low utilization. This way, we reduce the number of routes in those shifts.
(ii) Increasing the utilization of the shifts that (i) have a low workload or (ii) are less costly.

Experiments

We define five interventions, which we use in our experiments to improve the tactical routing schedule:

(i) Execute either heuristic A, B or both.
(ii) Prevent the balance of number of routes per shift to get worse.
(iii) Force the RVRP solver to improve the spread based on number of routes per shift.
(iv) Force the RVRP solver to reduce the fleet size.
(v) Allow consecutive deliveries to be within 18 hours.

Based on the best results from the first intervention (i), we experiment with applying interventions (ii), (iii) and (iv). The best result is used to improve further by applying intervention (v).

We found that when focusing on costs first, the best solution is found by applying heuristic A and B. Focusing on balancing workload first, the best solution found is by applying only heuristic B. Furthermore, we found that interventions (iv) contributed significantly in reducing the costs. Intervention (ii) and (iv) contributed significantly in improving the workload balance (i.e., they reduced the workload penalty WP_e).

Results tactical routing schedule

Using a hypothesis test, we found that reducing costs and improving workload balancing is not a tradeoff, but they are significantly positively related. The largest influence on this effect is intervention (iv) in which the RVRP solver is forced to reduce the fleet size. By reducing the fleet size, costs are directly reduced by €2000 per week due to lease costs, insurances and maintenance related to aging. Furthermore, when reducing the fleet size, automatically the largest shift(s) has also to be reduced in size, which improves the workload balancing. When the direct costs of €2000 per week are excluded from the model, we do not find a significant relation between fleet size and costs.

The best tactical routing schedule is the schedule that resulted from experiment XI in which we forced the RVRP solver to reduce the fleet size with 2 vehicles from 22 to 20 vehicles for the total network. With this tactical routing schedule, weekly costs are reduced by €5007. Additionally, €967 is saved by applying intervention (v). Workload balancing is measured using a penalty stating a normalized value of the sum of the squared deviations from the required workload balance. In our new tactical routing schedule, we reduced this penalty for the total network from 234,0 to 214,9.

Results operational policy

We performed a sensitivity analysis to test how the new routing schedule performs with deviating volumes. We find the new routing schedule being more sensitive to volume deviations. The savings are lower than the savings achieved in the tactical routing schedule that was based on forecasted volumes. The operational savings vary from €705 in a week with +8,0% volume, and €5406 in a week with -8,2% volume. The main reason is the impact of the fleet size. With the highest volume, the fleet size in the new tactical routing schedule is equal to the fleet size in the current situation. When volume decreases, in the new tactical routing schedule, we manage to reduce the fleet size, in contrast to the current situation, in which the fleet size remains the same as with the higher volume.

There are three measures, stating that the new routing schedule is less stable compared to the current situation. First, we found that the new routing schedule has 238 deliveries that changed more than 30 minutes from their originally planned time in the tactical routing schedule in a period of three weeks, compared to 168 changed delivery times in the current situation. Second, in the new routing schedule, in the period of three weeks, there are 30 changes in the number of routes per shift (i.e., a route being added or removed from a shift), compared to 2 changes in the current situation. Third, in the new tactical routing schedule, 18 stockouts occurred compared to 3 in the current situation However, the costs associated with the additional deliveries to replenish inventory are already included in the cost results shared above. The two explanatory results for this, are (i) the higher utilization of the trucks in the new routing schedule and (ii) the preference of Paragon to plan as early as possible. However, further research should be done regarding these issues, which we mention in Section 6.2.

Reflection on Core problems

In Section 1.5 we selected four out of nine core problems which we focus on:

1. The available inventory data is not sufficient to determine the actual storage capacity and construct feasible and efficient delivery patterns.

We achieved to determine the capacities of the customers' storage locations.

5. Wishes per restaurant are highly influencing the current delivery patterns, although HAVI should be in the lead to create more efficiencies.

The model we created, determines which delivery patterns are most efficient. HAVI can now take the lead in deciding on the delivery patterns. It has to be managed how to communicate this with the customers.

7. Having restaurants of multiple restaurant planners within one route, makes it cumbersome to discuss which volume (i.e., part of an order) must be pulled to preceding deliveries.

The operational policy we created, determines which volume must be delivered each delivery (in terms of HAVI, which volume is pulled to earlier deliveries). The restaurant planners have to adjust the orders such that they have the size proposed by our model.

8. The restaurant planners expect input from the transport department on which volume to pull, but do not get that input.

Instead of the transport department providing the input for the restaurant planners, the input now comes from our model.

6.2 Recommendations

In this section we provide recommendations for (i) practice and (ii) further research.

6.2.1 Recommendations for practice

We first give some recommendations that are a result of this research, on which we advise HAVI to implement our findings. We then give some recommendations, to improve the results and fully utilize the benefits from this research.

All the spreadsheets that are used for this research are based on the eastern part of the Netherlands. All spreadsheets should be extended to cover the total network. A benefit for this research is that the eastern part of the network does not consist any back-haulage customers. However, incorporating these customers, should be only minor changes to the spreadsheet, and they do not impact the logic of this research.

We advise HAVI to use the following recommendations by creating their new tactical routing schedules:

- Update the capacity values by updating the input data sources that are mentioned in Section 4.3. The capacities are updated by following Section 4.4.
- For all delivery units, determine the delivery windows by following Section 4.5.
- When it is known that the then active routing schedule is feasible, execute Section 4.6.
- Use the same settings as we used for experiment XI, because that delivered the best results.
    - While using the allocation procedure, use the priorities that have a primary focus on workload balancing.
    - Apply all nine steps of the algorithm of Section 4.7.2 (i.e., heuristic A and B)
    - When creating the routing schedule for the total network, intervention (iv) (i.e., reducing the fleet size) should be applied a minimum of two times, the first time a new tactical routing schedule will be created. It can be experimented if another execution of intervention (iv) will result in even better results. The following times the routing schedules will be created, again it should be experienced how often intervention (iv) can be executed. Maybe, when volumes increase due to seasonality, it is not even possible to reduce the fleet size, but it should be increased.

We advise HAVI to use the following recommendations in their daily business:

- Every day before the restaurant planners are starting to make orders, a system should run the algorithm of Section 4.9.2.
- Restaurant planners should adjust the orders such that the order quantities are the same as they are determined by Section 4.9.2.
- The transport planners should minimize the number of deliveries that are being delivered in another shift than they were planned in, in the tactical routing schedule.

Improvements

We advise HAVI to extent the historical data that is being used to determine lower bounds for the storage capacities. The more data is available, the better the lower bound that we can deduce.

We advise HAVI to develop a way of working in which the transport planners must look a few days in advance. The operational policy can be used to forecast the utilization of the routes. To help the restaurant planners to also act further in advance, we can use the operational policy that we created to guide them how to adjust the order quantities. Instead of the historical volumes that we used in this research to perform the sensitivity analysis, the operational policy will then be run with forecasted volumes. We then have an insight in what the utilization of the routes for the coming days will be after the utilization of the routes are being optimized. If, even after optimizing the utilization, the routes are planned above capacity, or way beneath capacity, the transport planners have an extra influence on the balancing of the workload. For example, if routes on a certain day are overloaded, but the restaurants have spare capacity, it is possible to insert an extra delivery on the prior day.

HAVI must decide whether they find the additional savings of intervention (v) enough to deliver consecutive deliveries within 18 hours apart from each other.

6.2.2 Recommendations for extended application

Thawing period

In Section 4.4 we determined the consumption period for each delivery unit. We assumed that goods can be consumed immediately after delivery. However, for the buns for example, they can only be consumed after a certain thawing period. We advise to research the impact of the assumption of direct consumption.

Product level

In Section 4.9.2 we developed the operational policy in which we determine per temperature zone how many delivery units should be delivered. The restaurant planners must make the orders manually by pulling products from consecutive deliveries to the first delivery. We advise to extend this model, so it states per product whether it must be delivered, and if so, how much it should be delivered. The translation of this model can be used to automate the process of making orders considering the utilization of the routes.

Routing of stockout deliveries

We assumed that stockouts are known sufficiently in time to be able to plan the extra deliveries within regular routes. No carriers are considered in this research. However, in practise, when a stockout takes place, sometimes, there is no time to wait until the next shift of departing routes. In those situations, carriers are hired. We advise to do further research about the impact of the carriers costs instead of the ability to plan the routes together with regular routes. If the impact is significant it can be considered to invest in a system or procedure that is able to more accurately predicts when a stockout is going to take place

6.2.3 Recommendations for further research

Approach comparison

In this research we used a revolutionary approach in which we jointly optimized the delivery patterns and delivery quantities. In our approach, the delivery patterns and delivery quantities are not decision variables but are a result from optimizing a routing schedule in which delivery units with corresponding delivery windows were optimized. We suggest further research in which this approach will be compared to existing heuristics for IRP models in which either (i) the delivery patterns and delivery quantities are the decision variables or (ii) each day will be determined which customers to deliver and which quantities to allocate. To have a fair comparison, at least the capacities of the customers' storage locations should be included. First, shelf life can be ignored, later it can be included. The comparison can be executed by applying the various approaches to multiple example instances, which are present in literature.

Stock out tracking

In Section 4.9.2 we stated that a stockout occurred when the inventory at the end of a day is below zero. However, two factors are not considered. Between the end of the day and the next delivery consumption can take place. Therefore, a stockout can take place, even if the inventory is not empty at the end of the day. Furthermore, every restaurant has a safety buffer in storage. This safety buffer is not considered in our research. Therefore, it is possible that we determined a stockout in this research, which in practice wouldn't be a stockout. We suggest further research can be done to include buffer storages into the model such that stockouts can be modelled more accurately. A sensitivity analysis can be executed to find the correlation between the buffer amount and the stockout occurrences.

Another variable that influences the stockout probabilities, is the volume that is used to make the tactical routing schedule. In our research the delivery windows are determined based on 110% of forecasted volumes. Increasing this percentage, will reduce the chance of stockouts, at the expense of smaller delivery windows.

Waste in the objective

In this research, the objective is to reduce costs and balance workload. Capacity and shelf life agreements are considered, but as restrictions. Because we create the tactical routing schedule based on fixed volumes, in practice, shelf life agreements can be violated when volumes deviate. Further research should be done to incorporate the waste reduction into a model in which delivery patterns and quantities are being optimized simultaneously.

References

Adulyasak, Y., Cordeau, J., & Jans, R. (2013). Formulations and Branch-and-Cut Algorithms for Multivehicle Production and Inventory Routing Problems. *INFORMS Jounal on Computing, 26*(1).

Andel, T. (1996). manage inventory, own information. *Transport and Distribution, 37*(5), 54-58.

Anily, S., & Federgruen, A. (1990). One Warehouse Multiple Retailer Systems with Vehicle Routing Costs. *Management Science, 36*(1), 92-114.

Archetti, C., Bertazzi, L., Hertz, A., & Speranza, M. (2012). A hybrid heuristic for an inventory-routing problem. *INFORMS Journal on Computing, 24*, 101-116.

Baptista, S., Oliveiro, R., & Zúquete, E. (2001). A period vehicle routing case study. *European Journal of Operational Research*(139), 220-229.

Bard, J., Huang, L., Jaillet, P., & Dror, M. (1998). decomposition approach to the inventory routing problem with satellite facilities. *Transportation Science, 32*, 189-203.

Bell, W., Dalberto, L., Fisher, M., Greenfield, A., Jaikumar, R., Kedia, P., . . . Prutzman, P. (1983). Improving the Distribution of Industrial Gases with an On-line Computerized Routing and Scheduling Optimizer. *Interfaces, 13*(6), 4-23.

Beltrami, E., & Bodin, L. (1974). Networks and vehicle routing for municipal waste collection. *Networks, 4*(1), 65-94.

Bertazzi, L., & Speranza, M. (2012). Inventory routing problems: an introduction. *EURO Journal on Transportation and Logistics*(1), 307-326.

Bertazzi, L., & Speranza, M. (2013). Inventory routing problems with multiple customers. *EURO Journal on Transportation and Logistics*(2), 255-275.

Bertazzi, L., Paletta, G., & Speranza, M. (2002). Deterministic Order-Up-To Level Policies in an Inventory Routing Problem. *Transportation Science, 36*(1), 119-132.

Bertazzi, L., Savelsbergh, M., & Speranza, M. (2008). Inventory Routing. In B. Golden, S. Raghaven, & E. Wasil, *The Vehicle Routing Problem: Latest Advances and New Challenges.* New York: Springer Science+Business Mdeia, LLC.

Campbell, A., & Savelsbergh, M. (2004a). Delivery volume optimization. *Transportation Science, 38*(2), 210-223.

Campbell, A., & Savelsbergh, M. (2004b). A Decomposition Approach for the Inventory-Routing Problem. *Transportation Science, 38*(4), 488-502.

Campbell, A., & Savelsbergh, M. (2004c). Efficient insertion heuristics for vehicle routing and scheduling problems. *Transportation Science, 38*(3), 269-378.

Campbell, A., & Wilson, J. (2014). Forty Years of Periodic Vehicle Routing. *Networks*, 2-15.

Carotenuto, P., Giordani, S., Massari, S., & Vagaggini, F. (2015). Periodic capacitated vehicle routing for retail distributions of fuel oils. *Elsevier: Transportation Research Procedia*(10), 735-744.

Coelho, L., Cordeau, J., & Laporte, G. (2012). Consistency in multi-vehicle inventory-routing. *Transportation Reseacrh Part C, 24*, 270-287.

Cordeau, J. F., Laporte, G., Savelsbergh, M., & Vigo, D. (2006). Vehicle Routing. In C. Barnhart, & G. Laporte, *Handbook in OR & MS, Vol. 14* (pp. 367-428). North Holland.

Dantzig, G., & Ramser, J. (1959). The Truck Dispatching Problem. *Management Science, 6*(1), 80-91.

Dror, M., & Ball, M. (1987). Inventory/routing: reduction from an annual to a short period problem. *Naval Research Logistics*(34), 891-905.

Dror, M., Ball, M., & Golden, B. (1985). Computational comparisons of algorithms for the inventory routing problem. *Annals of Operations Research*(4), 3-23.

Fisher, M., Greenfield, A., Jaikumar, R., & Kedia, P. (1982). *Real-time scheduling of a bulk delivery fleet: practical application of Lagrangian relaxation.* Philadelphia, U.S.A.: University of Pennsylvania.

Francis, P., Smilowitz, K., & Tzur, M. (2008). The Period Vehicle Routing Problem and its Extensions. In B. Golden, S. Raghavan, & E. Wasil, *The Vehicle Routing Problem* (pp. 73-102). Springer.

Golden, B., Raghavan, S., & Wasil, E. (2008). *The Vehicle Routing Problem Latest Advances and New Challenges.* Springer.

Hulshof, P. (2008). *How VMI Can Be Successful in Gas Distribution.* (Master's thesis). Retrieved from http://essay.utwente.nl.

Lahyani, R., Khemakhem, M., & Semet, F. (2014). Rich vehicle routing problems: From a taxonomy to a definition. *European Journal of Operational Research, 241*, 1-14.

Laporte, G. (2009). Fifty Years of Vehicle Routing. *Transportation Science, 43*(4), 408-416.

Laporte, G., Toth, P., & Vigo, D. (2013). Vehicle routing: historical perspective and recent contributions. *EURO Journal on Transportation and Logistics*(2), 1-4.

Lee, C., Bozer, Y., & White III, C. (2003). *A heuristic approach and properties of optimal solutions to the dynamic inventory routing problem.* Toronto, Ontario, Canada: University of Toronto.

Lei, L., Liu, S., Ruszczynski, A., & Park, S. (2006). On the Integrated Production, Inventory and Distribution Routing Problem. *IIE Transactions (Institute of Industrial Engineers), 38*(11), 955-970.

Lenstra, J., & Rinnooy Kan, A. (11). Complexity of vehicle routing and scheduling problems. *Networks, 11*(2), 221-227.

Liu, S. (2006). On the Integrated Production, Inventory and Distribution Routing Problem. *IIE Transactions, 38*(11), 955-970.

Moin, N., & Salhi, S. (2007). Inventory Routing Problems: A Logistical Overview. *The Journal of the Operational Research Society, 58*(9), 1185-1194.

Potter, A., Towill, D., & Disney, S. M. (2007). Integrating Transport into Supply Chains: Vendor Managed Inventory (VMI). In H. Jung, F. Chen, & B. Jeong, *Trends in Supply Chain Design and Management: Technologies and Methodologies* (pp. 331-344). London: Springer-Verlag.

Russell, W., & Igo, W. (1979). An assignment routing problem. *Networks*(9), 1-17.

Song, J., & Savelsbergh, M. (2007). Performance Measurement for Inventory Routing. *Transportation Science, 41*(1), 44-54.

Soysal, M., Bloemhof-Ruwaard, J., Haijema, R., & van der Vorst, J. (2015). Modeling an Inventory Routing Problem for perishable products with environmental considerations and demand uncertainty. *International Journal of Production Economics*(164), 118-133.

Trudeau, P., & Dror, M. (1992). Stochastic inventory routing: Route design sith stockouts and route failures. *Transportation Science, 26*, 171-184.

Waller, M., Johnson, M., & Davis, T. (1999). Vendor-managed inventory in the retail supply chain. *Journal of Business Logistics, 20*, 183-203.

Appendix I

Time windows

The time windows are already filled in within the master data of Paragon. The time windows are constructed in accordance with the transport planners and restaurant planners. The factors influencing the time windows are the following:

- Presence of employees (often one hour before opening).
- Legislation (e.g., city centers only early in the morning).
- Severe hinder for the operations or driver.
- Severe hinder for the customer.

The time windows are set such that times are excluded for which a penalty is risked, or operations of the customer or the driver are hindered severely. An example of time windows of McDonald's 'Laakhaven' is shown in Table 22.

Table 22: Example time windows McDonald's 'Laakhaven'

Day	Time windows
Monday	07:00-10:00 14:00-19:45
Tuesday	07:00-10:00 14:00-19:45
Wednesday	07:00-10:00 14:00-19:45
Thursday	07:00-10:00 14:00-19:45
Friday	07:00-10:00 14:00-19:45
Saturday	07:00-19:45
Sunday	09:00-19:45

We see that from Monday to Friday, the available delivery times are between 7 am and 10 am and between 2 pm and 7:45 pm. Between 10 am and 2 pm it is busy in this restaurant because it is located near a high school and the students get their lunch at the McDonald's. The weekend days have other time windows because there are no students studying those days. Like this restaurant, every restaurant has its own time windows, specific to its characteristics.

Appendix II

Center of Gravity

From different interviews with the operations manager, the manager of Restaurant Consulting McDonald's and a restaurant planner, it follows that there are various reasons why the workload is not always balanced to maximum potential. We must get an understanding why workload balancing must be dealt with, also when a better routing schedule is in place. For our example we assume the following:

- Deliveries takes place at Monday, Wednesday, Friday and Sunday.
- Delivery takes place before any consumption of that day.
- We deliver volume according to JIT.
- Demand is constant at 5 delivery units per day.

We use the term 'center of gravity' to have a simple measure of the volume balancing quantified in a single number. The CoG is important because the demand increases during the week. When HAVI would deliver according to JIT, there would be even greater peaks in the weekend.

In our example, assuming a demand of 5 delivery units per day, this results in $CoG = \frac{1}{35} * (1 * 10 + 3 * 10 + 5 * 10 + 7 * 5)$, which equals ${}^{125}/_{35}$. Now consider the change of Friday delivery to Thursday delivery, while the Monday and Sunday deliveries remain unchanged in volume. The impact on the storage Utilization is shown in Figure 17.

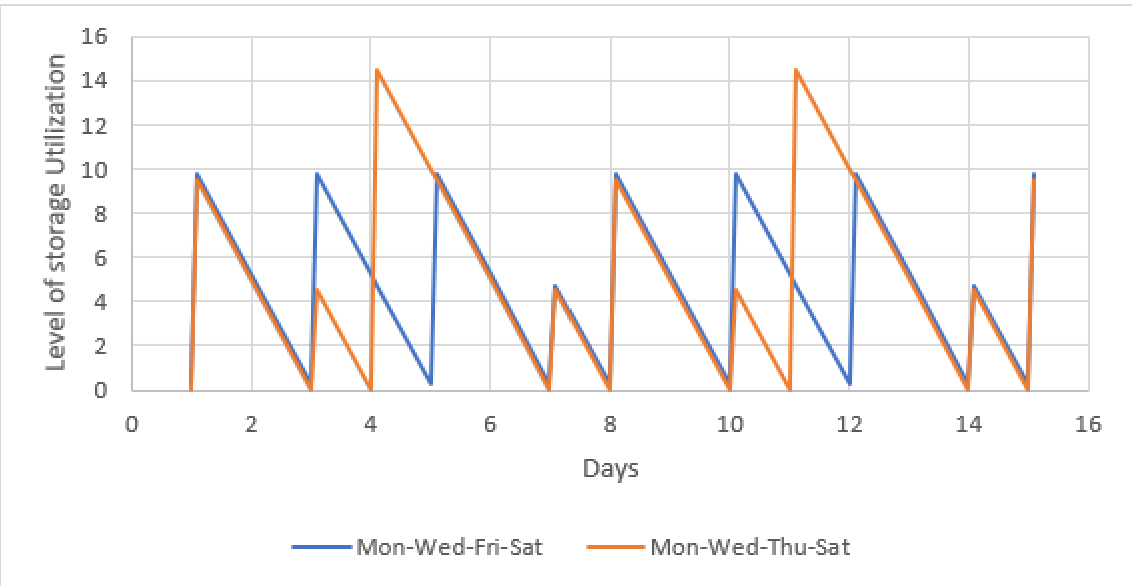

Figure 17: Storage levels for different delivery patterns

This results in a center of gravity of $CoG = \frac{1}{35} * (1 * 10 + 3 * 5 + 4 * 15 + 7 * 5)$, which equals ${}^{120}/_{35}$. Changing the delivery pattern has lowered the CoG. Changing a delivery to an earlier day, does not always decrease the CoG. For example, starting with a delivery pattern of Monday, Thursday, Friday and Sunday and changing the Thursday delivery to a Wednesday delivery increases the CoG from ${}^{120}/_{35}$ to ${}^{125}/_{35}$.

In Figure 17, we assumed JIT delivery. Consider the same delivery pattern as the blue line, but now we do not deliver according to JIT, but deliver some delivery units earlier than needed. The order quantities for the different ordering policies are as stated in Table 23. The impact on the storage levels is shown in Figure 18.

Table 23: Center of Gravity for different ordering policies

Ordering policy	Monday	Wednesday	Friday	Sunday	CoG
JIT	10	10	10	5	125/7
Earlier	13	12	7	3	105/7

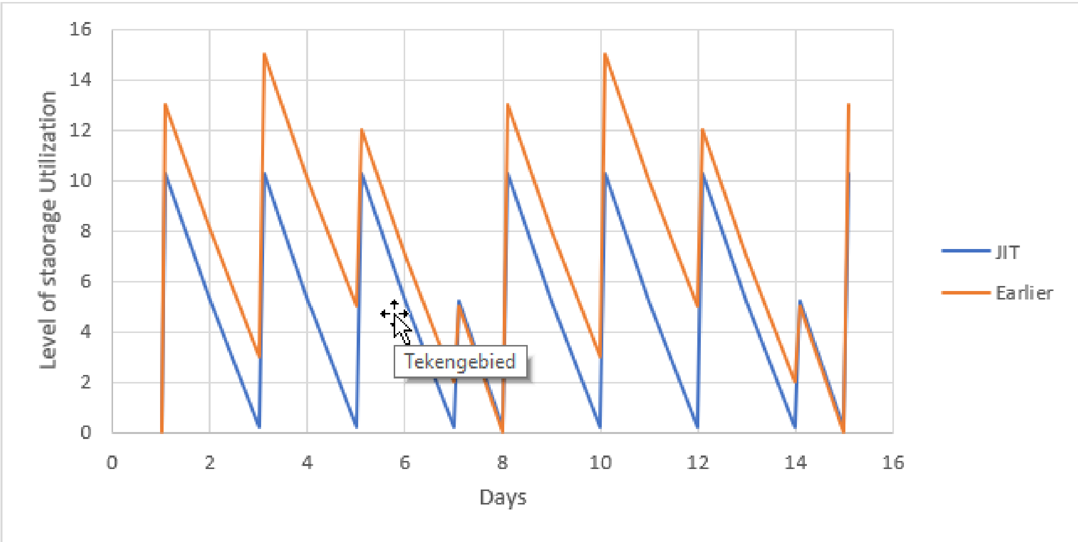

Figure 18: Storage levels for different ordering policies

Looking again at Figure 17, the blue pattern has a lower CoG. However, Pattern 2 gives the opportunity to pull some demand from the Wednesday delivery to the Monday delivery resulting in an even better CoG as we see in Table 23. In general, the CoG can always be improved by setting the first delivery in the week as early as possible. However, when applied to all restaurants, this will result in a peak on Monday, which is unpreferable. Concluding, there are two main factors influencing the CoG.

1. Changing the delivery pattern.
2. Changing the ordering policy (e.g., JIT).

In this research, we develop a method in which both the delivery patterns as the ordering policy is a decision variable. The CoG is measure that is important for the customers' perspective, because a lower CoG means that the average storage levels in the beginning of the week are higher.

Appendix III

JIT vs. Delivery balance

Because HAVI uses a system that proposes the orders for McDonald's based on JIT delivery, the restaurant planners must adjust those orders such that volume delivered is better balanced than volume consumed. These balances are shown in Figure 3. HAVI uses a dashboard that gives insight in the expected utilization of the shifts of the coming days when JIT delivery is applied. The planned utilization is shown in Figure 19 represented by the blue bars. We see that the routes on Sunday in the day shift, are planned above capacity. The restaurant planners should pull demand to earlier deliveries, resulting in the grey bars. Overloading the routes on the largest shifts is a tool for HAVI to force the restaurant planners to pull demand to earlier deliveries, such that delivered volume is better balanced. This is necessary because the system itself uses JIT delivery to propose the initial orders.

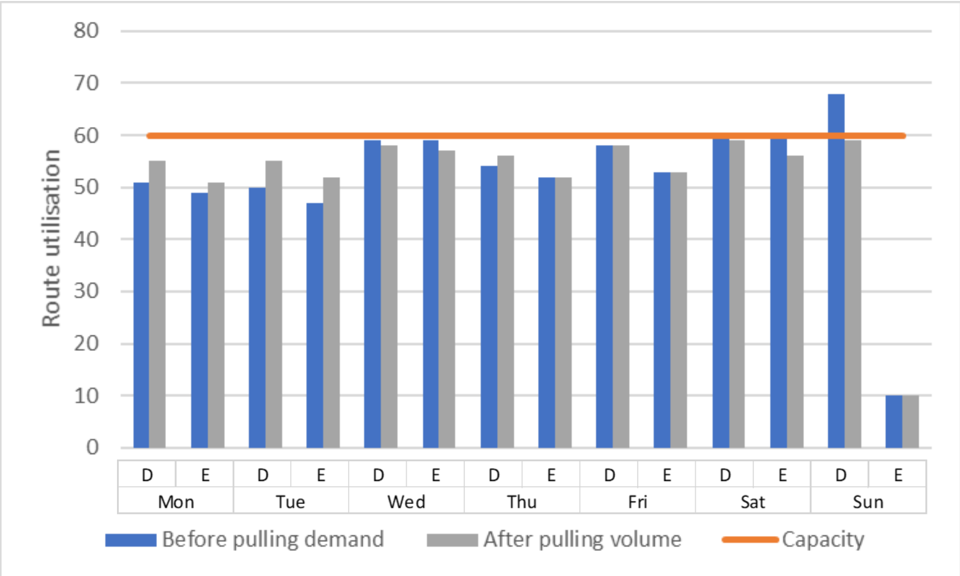

Figure 19: Rolling forecast utilization routes

Appendix IV

Examples from literature

Dantzig and Ramser (1959): *First to describe VRP*

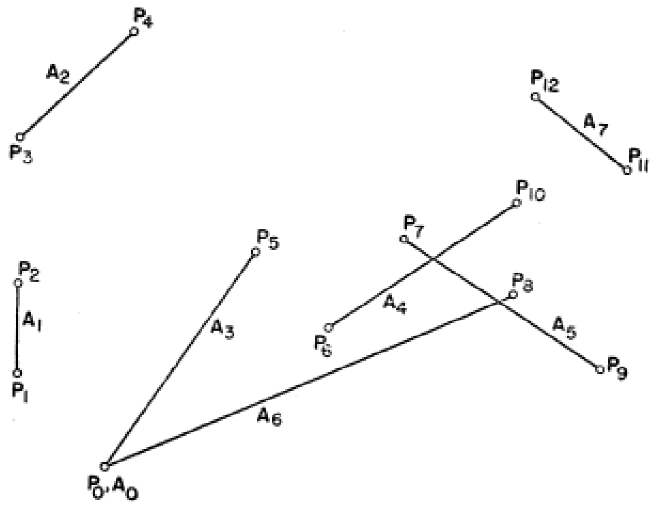

Figure 20: Pairing phase Dantzig and Ramser (1959)

The method of Dantzig and Ramser (1959) starts with carrying out sub optimizations by making pairs of delivery points. This phase is illustrated in Figure 20. Making a pair from P_3 to P_4, means that we include the trip from P_3 to P_4 into our routing schedule. The capacity of the truck is divided by equal parts. The size of the parts depends on the number of stages in the method. With two stages for example, at most half of the capacity of a truck can be used in the first stage making pairs. In the second stage the pairs are combined to sum up to at most the full capacity of the truck. The initial part size is set as follows: $\frac{C}{2^s}$ in which C is the capacity of a truck and s is the number of stages. The number of stages is dependent on the maximum number of locations within a route and is determined as follows: $s = RoundUp(\frac{Ln(L^{max})}{Ln(2)})$, in which L^{max} is the maximum number of locations within a route. It is also possible that a point is paired to the terminal (i.e., the DC in our case), more precise, even multiple points can be paired to the terminal point. The pairs are chosen such that the total distance to travel all pairs is minimized. In the second stage, the pairs made in the first stage, are combined in such a way that the total demand of points will not exceed the trucks capacity. Some extra procedures are included to improve the best solution found. For example, at each stage it is allowed that a pair is included only fractional. When such a situation occurs, we must determine which fractional pairs are placed in the best solution and which not, based on minimizing a sub problem. Based on this decision, the stage is recomputed with the pairs that we determined to be in the best solution being fixed, and the rest of the network as it was at the beginning of the stage. It is conjectured that the best solution found approaches the true optimal solution when the number of stations to be delivered increases and when the demand per station does not differ too widely.

Bell et al. (1983): *First to describe IRP*

In an IRP, at each stage a decision is made, which customers to deliver and what quantities to deliver. This problem is first described by Bell et al. (1983). They used a simple example to illustrate the benefits of dealing with the routing problem whilst the decision maker can make the route and simultaneously determine the delivery days and delivery quantities. The problem is illustrated in Figure 21. For four customers a periodic schedule must be determined which fulfils the demand of the customers.

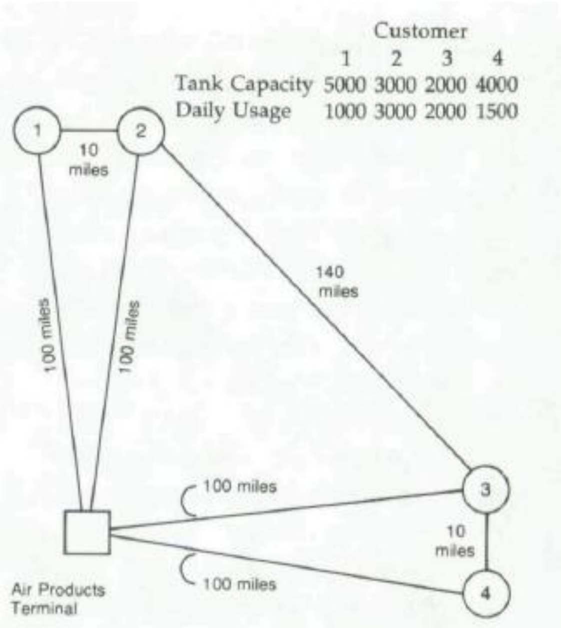

Figure 21: Simple example IRP from Bell et al. (1983)

The trivial solution would be to deliver customer 1 and 2 in one route every day with a total volume of 4000, and customer 3 and 4 in one route every day with a total volume of 3500. Every day there are two routes of each 210 miles. The average mileage per day is 420 miles. A better solution can be found when making a periodic schedule of two days. On the first day, customer 2 and 3 are being delivered with quantities of 3000 and 2000, respectively. On the second day two routes are made. These routes are the same as in the trivial solution. The first day 340 miles are driven, and the second day again 420 miles. The average mileage is decreased from 420 miles to 380 miles per day. The customers' demand is still fulfilled, and the average delivery frequency is decreased (customer 1 and customer 4 are delivered once per two days instead of every day).

Appendix V

Paragon as RVRP solver

In this research, we use Paragon as RVRP solver because HAVI already uses Paragon in creating their tactical routing schedules. An RVRP solver is a routing optimization software that is able to solve VRP problem with a lot of practical constraints. In this appendix the characteristics of Paragon are discussed.

The objective within Paragon can be chosen. The options are as follows:

- Minimize duty time
- Minimize distance
- Minimize duty time (clustered)
- Minimize distance (clustered)

Minimizing distance is straightforward. Minimizing duty time is the minimization of driving time, waiting time, loading and unloading time and an extra duty time that corresponds to the administration of a route. This extra duty time is set at one and a half hour.

The other two options do the same, but then clustered. Which means, that the distance and driving time, from and to the DC is excluded. This results in a routing schedule in which most routes consist of customer that are closer to each other than with the other objectives. The experience is that the objective with clustering results in a routing schedule in which it is easier to adjust routes when some routes are above or way below capacity. Therefore, in this research, when we say that we optimize using the RVRP solver, we use Paragon the minimize the duty time (clustered). Other characteristics of Paragon are:

- Paragon considers asymmetric driving times and distances, that are time dependent.
- For some customers there is a penalty incorporated when the customer is visited within some time boundaries.
- Paragon has loading times that are dependent on the delivered quantity. The delivery time is always a fixed time, plus a variable time per delivery unit.
- Paragon uses depot smoothing, which means that within some time intervals there is a maximum number of trucks that can depart from the depot. This is in place to balance the workload of loaders.

Appendix VI

Example improving by shifting partial deliveries

Looking at a simplistic example, we show the benefits of our new approach. Consider Figure 22 showing six customers with their initial assignment. Customer 1 to 5 have a fixed delivery day. The goods for customer 6 are split in two parts: The first part should be delivered on Monday or Tuesday; the second part should be delivered on Tuesday or Wednesday. In the initial solution, every customer gets one delivery. Every day, one route is driven. The number next to the arrow is the duration to travel from one to the other customer. The initial solution has a driving time of 12 hours and 45 minutes. Assuming a delivery time of 30 minutes per stop, the total duration is 15 hours and 45 minutes.

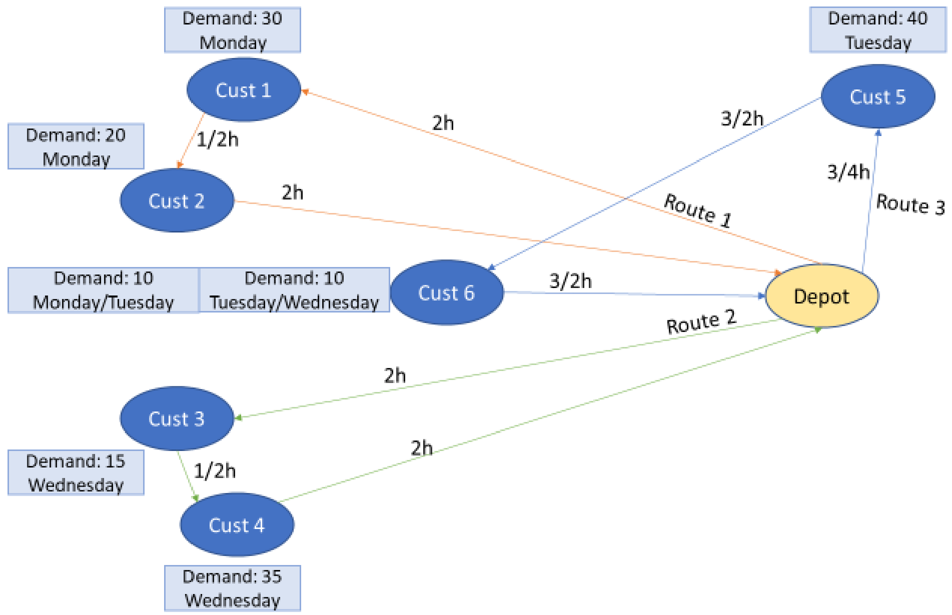

Figure 22: Initial solution simplistic example

In a traditional approach, we cannot switch the volume of a customer to another delivery day in the improvement stage. Often the delivery days are being determined in a previous phase. Because we now determined for every delivery unit a delivery window, it is possible to find an improvement. The demand of customer 6 is split over two routes on two different days. We show the improved solution in Figure 23. The improved solution has a driving time of 11 hours and 30 minutes. The total duration is 15 hours because we added 30 minutes per stop for seven stops (two for customer 6). The improved solution saves 45 minutes by adding a delivery moment for customer 6.

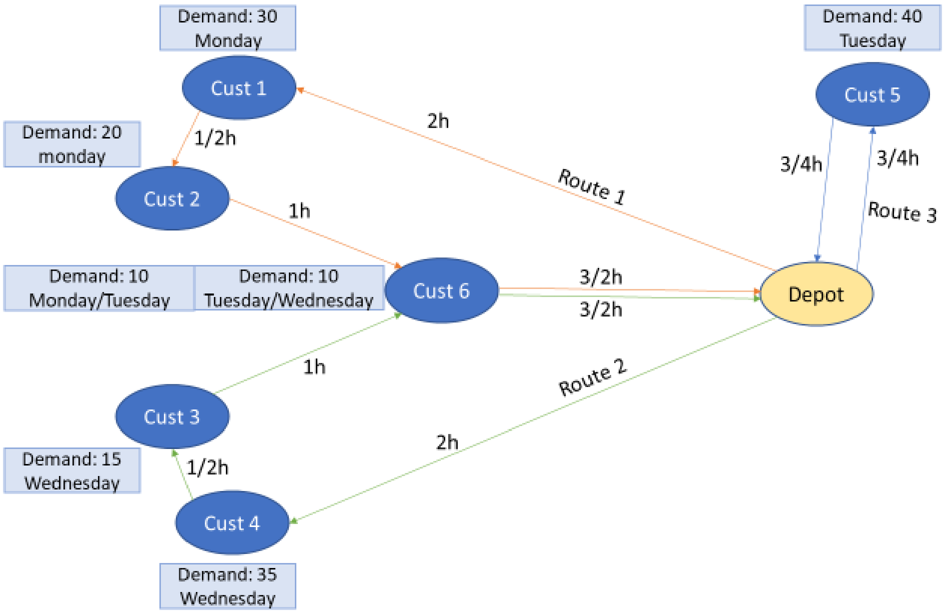

Figure 23: Improved solution simplistic example

Appendix VII

Example determining delivery windows

We look at an example of how delivery windows are determined. For this example, we use McDonald's 'Winterswijk'. To keep the example simple, assume the following: We only look at the fridge; the demand for chilled products is three roll containers per week and the capacity of the fridge is equal to two roll containers. This restaurant gets his deliveries at Monday, Wednesday, Friday, and Sunday. Every day the delivery takes place at 12 am. Expressed in hours in the week, the deliveries take place at hour 8, 56, 104, 152 resp. In Figure 24 we show the usage of the fridge over the span of a week. The blue line is the actual usage expressed in roll containers stored. The consumption per hour is based on the historical hourly sales distribution, with a total of three roll containers in the week. To determine the delivery windows, we first look at the limitations of the capacity.

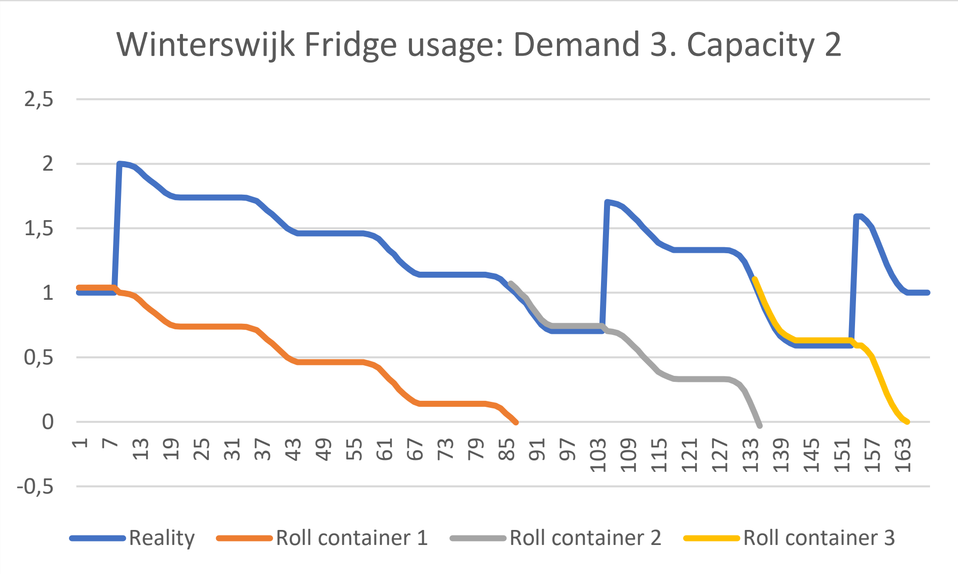

Figure 24: Fridge usage example

The initial storage is set at one roll container. The orange line shows the usage of the fridge in case that the first roll container is not yet delivered. The first roll container cannot be delivered after the 87th hour (i.e., Thursday 3 pm), because the usage of the fridge would have reached zero. When the first delivery unit would be delivered at the 87th hour, it would immediately be used for consumption. From Equation 5 it follows that $Latest(U_{C,T,n})$ as well as $t_{C,T,n}^{start}$ are being equal to 87.

Because the initial inventory is set at one roll container, the inventory level at the end of the previous week is also one roll container. The first roll container cannot be delivered earlier than Sunday after all sales, because the usage of the fridge would have been above capacity. From Equation 2 it follows that

$$Earliest(U_{Winterswijk,Chilled,1})_{Cap} = t^{end}_{Winterswijk,Chilled,(1-(Cap_{Winterswijk,Chilled}))} = t^{end}_{Winterswijk,Chilled,(-1)}.$$

Because we deal with a repetitive schedule, we take the last delivery unit of the previous week as delivery unit 0. Delivery unit -1 would the delivery unit before delivery unit 0, which is delivery unit 2 of the previous week. This makes $t^{end}_{Winterswijk,Chilled,(-1)} = t^{end}_{Winterswijk,Chilled,2}$.

Given that the first roll container is being delivered on Monday morning, the second roll container cannot be delivered on Wednesday Morning at the 56th hour. Because the usage of the fridge is still above one, and added with the roll container being delivered, the usage would be higher than capacity of two. The second roll container cannot be delivered earlier than the 87th hour (i.e., Thursday 7 pm), and not later than the 135th hour (i.e., Saturday 7 pm). The third roll container cannot be delivered earlier than the 135th hour, and not later than the 9th hour (i.e., Monday 1 pm following week). The second roll container is being consumed until Sunday before midnight. However, there are no sales between Sunday midnight and Monday 9 am (this is not true, but it then shows a feature of our method of determining the consumption period). Therefore, we can say that the end of the consumption period is Monday 9 am. We use this difference in our favor. We choose the end of a consumption period as early as possible and the beginning of a consumption period as late as possible. Now we found the delivery windows in the case of unlimited shelf life. We show the delivery windows without shelf life limitation in Table 24.

Table 24: Delivery windows without shelf life limitation

Roll Container	Earliest delivery moment	Latest delivery moment
1	Sunday end of day	Thursday 7 pm
2	Thursday 7 pm	Saturday 7 pm
3	Saturday 7 pm	Monday 1 pm

In this specific example, without limitation of the shelf life and a capacity of two roll containers, the consumption period of a roll container is equal to the delivery window of the next roll container (except for the cases in which there is no sales in the hours following the consumption period). In Figure 25 we show the Consumption periods in combination with the delivery periods for a period of two weeks. We see that the delivery periods are equal to the consumption period of the previous roll container (except for the period in which no goods are consumed). We see a gap between the consumption period of roll container 2 and 3. This is because no goods are consumed between Sunday night and Monday morning. In other words, a roll container cannot be delivered earlier than the end of the consumption period of the roll container two places earlier.

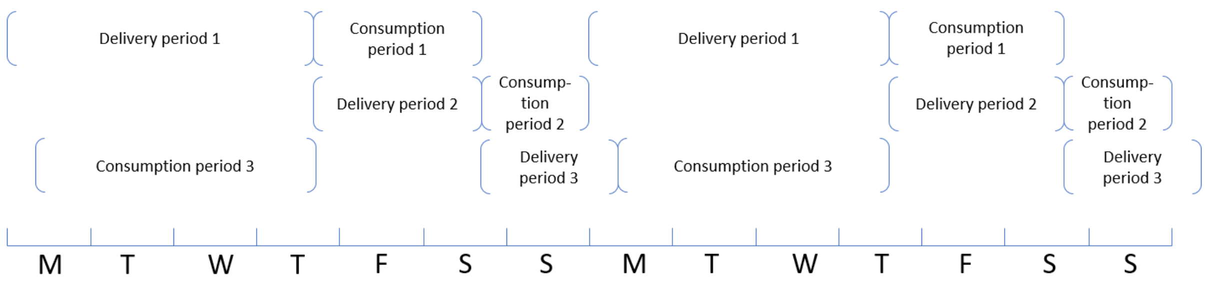

Figure 25: Consumption and delivery periods without shelf life limitation

Now, we continue by determining the delivery windows including the shelf life limitations. We know that the end of the consumption period of a roll container is equal to the end of the delivery window of the next roll container. The exception is that the delivery ends with a period of no sales, then the consumption period ends at the last moment of sales. The ends of the consumption periods of the roll containers are at hours 135, 168 and 87 resp. The 135th hour is on Saturday. The minimum shelf life agreements with the customer states that upon delivery, the goods should have a minimum shelf life left of four days (including the delivery day). Therefore, the first roll container cannot be delivered earlier than Wednesday (4 am). We already knew that the roll container could not be delivered earlier than Monday 4 am, but this changes to Wednesday 4 am. We show the impact of the shelf life limitation on the delivery period in Figure 26.

Continuing with the second roll container, the 168th hour is on Sunday (Monday before 4 am is counted as Sunday), which means that the second roll container cannot be delivered earlier than Thursday 4 am. We already knew that the roll container could not be delivered earlier than Thursday 3 pm, this remains the same, because this is already more restrictive than the shelf life limit. The 87th hour is on Thursday. Therefore, the third roll container cannot be delivered earlier than Monday 4 am. Now we see the benefits of choosing the beginning of the consumption period being Monday 9 am instead of

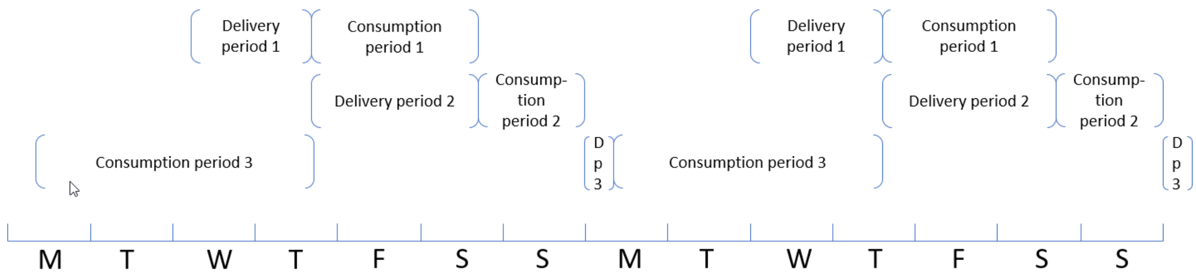

Figure 26: Consumption and delivery periods with shelf life limitation

Sunday before midnight. We show the delivery windows with shelf life limitation in Table 25. We see that two out of the three delivery windows are more restricted by adding the shelf life limitations. The delivery window gives us the boundaries wherein the delivery time is chosen. Therefore, the possible delivery moments of the first roll container, is part of the time windows of the specific customer (McDonald's 'Winterswijk') that lies between Wednesday 4 am and Thursday 3 pm.

Table 25: Delivery windows with shelf life limitation

Roll Container	Earliest delivery moment	Latest delivery moment
1	Wednesday 4 am	Thursday 3 pm
2	Thursday 3 pm	Saturday 3 pm
3	Monday 4 am	Monday 9 am

The first roll container must be delivered at the Wednesday delivery. The second roll container must be delivered on Friday. The third roll container must be delivered at the Monday delivery. It is not necessarily the case that a roll container only fits in one delivery moment. From this example we see that it is not the case that the first roll container is also being delivered first. The first roll container is the roll container that is consumed first in the week. Therefore, the first roll container of a week is delivered in the previous week.

Appendix VIII

Setting priorities and assigning points to the shifts

In our algorithm we use the allocation procedure several times. We distinguish between emphasizing on low priority shifts and high priority shifts. We also distinguish between either focusing on cost reduction first or focusing on balancing workload first. Therefore, four scenarios are defined.

Focus on costs first, emphasizing high priority shifts

In this scenario the points assigned to the shifts are as in Table 26. We always start with the allocating delivery units to the shift with the highest number of points. Delivery units with the least points are being allocated first. When allocating the delivery units to the Tuesday Evening shift, all delivery units that can be allocated have at least 100 points. The first delivery units that will be allocated are the ones that can only be delivered in a Tuesday evening route. Then the delivery units will be allocated that can be delivered on a Tuesday evening route as well as a Sunday evening route. This way, Sunday evening is avoided as much as possible. The more alternatives and the worse the alternatives are, the higher the chance that we allocate the delivery unit to the Tuesday evening. Assigning the points is quite arbitrary. We have for example a large gap between Tuesday daytime and Thursday daytime because adding workload to the shifts from Tuesday daytime and lower

Table 26: Points - Costs first, emphasize high priority

Priority	#1	#2	#3	#4	#5	#6	#7	#8	#9	#10	#11	#12	#13	#14
Shift	Tue E	Thu E	Wed D	Wed E	Fri E	Mon E	Thu D	Tue D	Fri D	Mon D	Sat D	Sat E	Sun D	Sun E
Points	100	90	80	70	60	50	40	15	12	7	4	3	2	1

Focus on costs first, emphasizing low priority shifts

Table 27: Points - Costs first, emphasize low priority

Priority	#1	#2	#3	#4	#5	#6	#7	#8	#9	#10	#11	#12	#13	#14
Shift	Tue E	Thu E	Wed D	Wed E	Fri E	Mon E	Thu D	Tue D	Fri D	Mon D	Sat D	Sat E	Sun D	Sun E
Points	1	2	4	5	9	11	15	40	50	75	100	150	200	250

Focus on balancing first, emphasizing high priority shifts

Table 28: Points - Balancing first, emphasize high priority

Priority	#1	#2	#3	#4	#5	#6	#7	#8	#9	#10	#11	#12	#13	#14
Shift	Sun E	Tue E	Thu E	Sat E	Wed E	Wed D	Sun D	Fri E	Mon E	Thu D	Tue D	Mon D	Sat D	Fri D
Points	100	90	80	70	60	50	40	20	12	7	4	3	2	1

Focus on balancing first, emphasizing low priority shifts

Table 29: Points - Balancing first, emphasize low priority

Priority	#1	#2	#3	#4	#5	#6	#7	#8	#9	#10	#11	#12	#13	#14
Shift	Sun E	Tue E	Thu E	Sat E	Wed E	Wed D	Sun D	Fri E	Mon E	Thu D	Tue D	Mon D	Sat D	Fri D
Points	1	2	3	5	7	10	15	20	25	50	80	100	120	250

Appendix IX

Interventions

Intervention (i) and (v) do not need further elaboration. We elaborate on intervention (ii), (iii) and (iv).

Intervention (ii)

For this intervention we determine the number of routes in the current situation in the East of the Netherlands. Also, we decide what the maximum routes allowed is for each shift. The number of routes per shift (V_s^{cur}) in the current situation is stated in Table 30.

Table 30: Number of routes in the current situation

Shift	Mon D	Mon E	Tue D	Tue E	Wed D	Wed E	Thu D	Thu E	Fri D	Fri E	Sat D	Sat E	Sun D	Sun E
V_{cur}^s	4	6	3	2	3	6	3	3	3	5	4	4	4	0

When determining the maximum number of routes (V_s^{max}) allowed for each shift, we distinguish between focusing on costs first, or focusing on balancing workload first. The priorities of the shifts for these scenarios are stated in Appendix VI. The maximum number of routes allowed for each shift are stated in Table 31 and Table 32.

Focusing on costs first

Table 31: Maximum number of routes per shift, when focusing on costs first

Shift	Mon D	Mon E	Tue D	Tue E	Wed D	Wed E	Thu D	Thu E	Fri D	Fri E	Sat D	Sat E	Sun D	Sun E
V_{max}^s	4	6	4	6	6	6	6	6	3	6	4	5	4	0

Focusing on balancing first

Table 32: Maximum number of routes per shift, when focusing on balancing first

Shift	Mon D	Mon E	Tue D	Tue E	Wed D	Wed E	Thu D	Thu E	Fri D	Fri E	Sat D	Sat E	Sun D	Sun E
V_{max}^s	4	6	3	2	3	6	3	3	3	5	4	4	4	6

Example of delivery window for dummy customer

The dummy customer is used as an occupier. A dummy customer routed in a shift decreases the possible number of normal routes in that shift with one. The delivery windows for the dummy customer are constructed such that they can only be delivered in shifts with equal or lower priority. When a dummy customer is transferred to a shift with lower priority, another route can be transferred from the lower priority to the higher priority shift. For the priorities of the shifts, see the example of Table 7. The dummy customer is located near to the DC. Therefore, the dummy customer is delivered within the time windows that departure of routes is possible. For a dummy customer initially delivered on Tuesday daytime for example, the delivery window is as follows:

Mon 04:00-08:30 & Tue 04:00-08:30 & Sat 04:00-08:30 13:00-18:00 & Sun 04:00-08:30 13:00-18:00

Intervention (iii)

For intervention (iii) we take the same V_s^{max} as in intervention (ii) as stated in Table 31 and Table 32. We distinguish between either focusing on costs first or focusing on balancing workload first. When focusing on costs first, the set of shifts from which a route must be transferred to a shift with higher priority is all shifts on which salary supplements apply. When focusing on balancing workload first, the set of shifts from which a route must be transferred to a shift with higher priority is all shifts except for the lowest priority (i.e., Sunday night).

We randomly select a delivery unit that is being delivered within the determined set of shifts. The whole delivery is removed from the routing schedule, thus also all other delivery units that are being delivered to the same customer in the same route. We then reoptimize the routing schedule and check whether the RVRP solver has saved a route. If a route is saved within shift with priority p that is within the determined set of shifts, we add a delivery to the dummy customer in that shift and we remove a delivery to the dummy customer from the shift with priority $p - 1$. For example, when focusing on costs after some delivery moments are being removed, the RVRP has saved a route in the Saturday evening shift. We add a delivery to the dummy customer on Saturday evening (priority 12 See Table 26) with the following time windows: Sat 13:00-18:00 & Sun 04:00-08:30 13:00-18:00. We also remove a delivery to the dummy customer. This delivery is removed from Saturday daytime (priority 11) with the following time windows: Sat 04:00-08:30 13:00-18:00 & Sun 04:00-08:30 13:00-18:00.

Intervention (iv)

In Section 4.7.4 we gave the method that can be applied for intervention (iv) applied on the total network. In this research, a part of the total network is in scope and we therefore need an adjusted procedure for intervention (iv). The procedure is as follows:

To reduce the fleet size, we must know what the maximum number of vehicles needed is in the current situation. As we see in Figure 14 in Section 4.7.1, a maximum of 22 vehicles are needed in the current situation. To reduce the fleet size with 1 vehicle, the maximum vehicles for the total network allowed is 21. We use the same variables as in Section 4.7.4, but add a superscript stating whether the variable applies to the total network (NL) or a part of the network (East).

- Set $V_s^{max,East} = V_s^{cur,East} + (V^{max,NL} - V_s^{cur,NL}) - 1$ for all shifts s for which $V_s^{cur,NL} < V^{max,NL}$
- Set $V_s^{max,East} = V_s^{cur,East} + (V^{max,NL} - V_s^{cur,NL})$ for all shifts s for which $V_s^{cur,NL} = V^{max,NL}$
- Remove deliveries from the set of shifts for which $V_s^{cur,NL} = V^{max,NL}$ until it is possible to save a route in each of these shifts with the RVRP solver.
- Set $V_s^{max,East} = V_s^{cur,East} + (V^{max,NL} - V_s^{cur,NL}) - 1$ for all shifts in which we saved a route.
- Reoptimize the routing schedule with the RVRP solver by routing all deliveries that we removed earlier.

Appendix X

Experimental results

In this appendix we state the explanatory experimental results. In Section 5.2, the main results are discussed. First, we mention the results for the tactical routing schedule. Thereafter, we mention the results of the operational policy.

Results tactical routing schedule

Table 33: Routes per shift

Experiment	Mon D	Mon E	Tue D	Tue E	Wed D	Wed E	Thu D	Thu E	Fri D	Fri E	Sat D	Sat E	Sun D	Sun E	Total
Base	4	6	3	2	3	6	3	3	3	5	4	4	4	0	50
I	4	6	4	2	3	6	3	3	3	5	4	4	5	1	53
II	6	4	3	2	3	6	4	2	3	5	3	5	4	0	50
III	4	6	3	2	3	6	3	3	3	5	4	4	5	0	51
IV	5	6	3	2	3	5	3	4	3	4	4	4	4	1	51
V	5	4	3	2	3	6	3	3	3	5	4	4	4	1	50
VI	6	4	3	2	3	5	3	4	2	5	4	4	5	0	50
VII	4	6	3	2	3	6	3	3	3	5	4	4	4	0	50
VIII	4	5	3	2	3	6	3	3	3	5	3	4	4	0	48
IX	5	4	3	2	3	5	3	5	3	6	3	4	3	0	49
X	4	6	3	2	3	6	3	3	2	5	4	4	4	0	49
XI	3	6	2	3	3	6	3	4	1	6	3	4	4	0	48
XII	4	6	3	2	3	6	3	3	3	5	4	4	4	1	51
XIII	4	4	3	2	3	6	3	3	3	5	4	4	5	0	49
XIV	4	3	3	1	4	6	2	4	3	5	3	5	6	0	49
XV	4	6	3	2	3	6	3	3	2	5	4	4	5	0	50
XVI	3	5	2	3	4	5	3	4	1	5	3	6	5	0	49
XVII	3	6	2	2	4	4	3	4	1	6	3	5	4	0	47

Table 34: Delivery units per shift and warehouse related costs

Experiment	200%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	150%	150%	200%		€ 4
Base	**Mon D**	**Mon E**	**Tue D**	**Tue E**	**Wed D**	**Wed E**	**Thu D**	**Thu E**	**Fri D**	**Fri E**	**Sat D**	**Sat E**	**Sun D**	**Sun E**	**Total**	**Warehouse €**
I	183	284	164	117	162	293	175	137	151	279	192	229	201	0	2567	€ 11.860
II	239	346	153	25	104	273	138	97	162	296	231	237	253	25	2579	€ 12.352
III	264	161	156	109	167	332	223	117	133	242	166	258	239	0	2567	€ 12.318
IV	164	291	159	97	164	310	174	177	119	270	223	205	214	0	2567	€ 11.762
V	265	314	163	78	127	237	164	153	156	233	233	217	235	6	2581	€ 12.312
VI	246	161	131	114	175	299	174	163	124	274	219	240	224	23	2567	€ 12.272
VII	229	212	149	78	135	264	170	219	101	252	233	238	287	0	2567	€ 12.234
VIII	182	268	146	94	143	318	166	167	162	293	215	196	223	0	2573	€ 11.858
IX	186	272	172	71	143	305	169	170	170	280	179	226	230	0	2573	€ 11.948
X	257	226	149	86	155	272	173	245	105	341	170	214	180	0	2573	€ 12.108
XI	215	266	153	75	150	296	180	174	116	294	230	197	227	0	2573	€ 12.000
XII	175	307	115	111	135	313	153	198	60	350	180	236	240	0	2573	€ 11.944
XIII	150	237	158	120	168	333	176	173	110	256	199	240	221	26	2567	€ 11.894
XIV	214	182	141	105	152	312	167	178	146	236	227	228	279	0	2567	€ 12.138
XV	220	177	177	58	190	294	117	200	141	249	174	249	321	0	2567	€ 12.288
XVI	204	178	160	115	173	329	174	172	100	253	224	223	262	0	2567	€ 12.054
XVII	176	213	119	123	225	263	179	180	51	260	173	323	282	0	2567	€ 12.182
	166	297	118	109	218	211	167	221	38	350	168	274	236	0	2573	€ 11.976

Table 35: Utilization per shift

Experiment	Mon D	Mon E	Tue D	Tue E	Wed D	Wed E	Thu D	Thu E	Fri D	Fri E	Sat D	Sat E	Sun D	Sun E	Average
Base	76%	79%	91%	98%	90%	81%	97%	76%	84%	93%	80%	95%	84%		86%
I	100%	96%	64%	21%	58%	76%	77%	54%	90%	99%	96%	99%	84%	42%	81%
II	73%	67%	87%	91%	93%	92%	93%	98%	74%	81%	92%	86%	100%		86%
III	68%	81%	88%	81%	91%	86%	97%	98%	66%	90%	93%	85%	71%		84%
IV	88%	87%	91%	65%	71%	79%	91%	64%	87%	97%	97%	90%	98%	10%	84%
V	82%	67%	73%	95%	97%	83%	97%	91%	69%	91%	91%	100%	93%	38%	86%
VI	64%	88%	83%	65%	75%	88%	94%	91%	84%	84%	97%	99%	96%		86%
VII	76%	74%	81%	78%	79%	88%	92%	93%	90%	98%	90%	82%	93%		86%
VIII	78%	91%	96%	59%	79%	85%	94%	94%	94%	93%	99%	94%	96%		89%
IX	86%	94%	83%	72%	86%	91%	96%	82%	58%	95%	94%	89%	100%		88%
X	90%	74%	85%	63%	83%	82%	100%	97%	97%	98%	96%	82%	95%		88%
XI	97%	85%	96%	62%	75%	87%	85%	83%	100%	97%	100%	98%	100%		89%
XII	63%	66%	88%	100%	93%	93%	98%	96%	61%	85%	83%	100%	92%	43%	84%
XIII	89%	76%	78%	88%	84%	87%	93%	99%	81%	79%	95%	95%	93%		87%
XIV	92%	98%	98%	97%	79%	82%	98%	83%	78%	83%	97%	83%	89%		87%
XV	85%	49%	89%	96%	96%	91%	97%	96%	83%	84%	93%	93%	87%		86%
XVI	98%	71%	99%	68%	94%	88%	99%	75%	85%	87%	96%	90%	94%		87%
XVII	92%	83%	98%	91%	91%	88%	93%	92%	63%	97%	93%	91%	98%		91%

Table 36: Total duty time per shift and corresponding costs

Experiment	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	150%	150%	200%	200%		€ 30
Shift	**Mon D**	**Mon E**	**Tue D**	**Tue E**	**Wed D**	**Wed E**	**Thu D**	**Thu E**	**Fri D**	**Fri E**	**Sat D**	**Sat E**	**Sun D**	**Sun E**	**Total**	**Driver €**
Base	35,3	45,5	32,1	16,7	27,7	43,5	28,0	22,9	30,2	40,3	30,4	35,2	31,7	0,0	419,517	€ 14.520
I	37,0	46,3	36,6	9,3	24,6	43,8	26,7	20,5	28,3	42,3	37,4	33,9	33,5	7,2	427,217	€ 15.106
II	48,7	28,5	32,1	14,6	26,8	45,3	40,1	17,9	28,1	39,4	25,2	47,4	31,5	0,0	425,567	€ 14.801
III	31,1	45,6	30,5	11,8	26,5	44,5	27,2	24,7	27,1	41,5	36,6	32,1	30,6	0,0	409,55	€ 14.233
IV	42,3	45,6	30,0	12,4	24,3	38,4	28,5	24,3	28,3	34,9	37,1	34,9	32,5	3,1	416,567	€ 14.644
V	40,8	27,8	26,8	14,0	24,8	43,2	27,1	21,9	27,3	41,4	37,7	35,5	30,4	9,0	407,667	€ 14.508
VI	42,7	31,4	27,6	11,1	21,5	40,3	27,5	25,8	24,9	40,4	33,4	32,8	40,6	0,0	399,717	€ 14.201
VII	31,6	42,4	28,6	13,4	21,2	47,3	30,0	22,7	28,6	42,1	35,9	31,7	30,4	0,0	405,8	€ 14.100
VIII	31,9	40,4	30,7	11,2	20,7	45,9	27,9	23,2	29,8	41,8	28,3	33,4	31,2	0,0	396,417	€ 13.755
IX	43,6	31,0	28,1	12,8	21,2	40,3	30,7	39,5	24,7	52,7	27,6	33,9	30,5	0,0	416,433	€ 14.330
X	33,1	42,2	28,9	11,7	24,2	44,1	32,8	23,4	26,0	43,0	37,5	31,8	30,5	0,0	409	€ 14.224
XI	25,5	47,2	22,8	18,0	21,1	44,8	27,4	30,4	10,0	58,4	25,9	35,6	33,7	0,0	400,817	€ 13.959
XII	32,3	45,0	30,8	15,6	25,8	47,0	27,7	23,3	27,9	39,6	32,4	33,6	32,4	4,8	418,117	€ 14.650
XIII	37,7	30,8	27,3	14,8	23,7	44,0	26,2	24,6	29,6	37,4	32,6	32,1	39,1	0,0	399,75	€ 14.136
XIV	36,9	29,3	30,4	8,0	30,5	40,7	17,5	31,9	26,6	39,8	24,9	39,9	44,5	0,0	400,817	€ 14.333
XV	34,8	36,3	29,4	13,9	26,3	45,7	27,7	23,2	23,6	41,4	37,1	32,3	34,3	0,0	405,9	€ 14.246
XVI	31,2	37,9	20,4	19,9	33,6	35,4	27,5	29,1	12,3	44,3	23,7	46,8	37,7	0,0	399,633	€ 14.176
XVII	26,6	46,7	23,3	16,3	34,1	28,9	26,5	32,2	9,3	51,5	26,6	41,1	31,6	0,0	394,383	€ 13.794

Table 37: Total traveled distance per shift and corresponding costs

Experiment	Mon D	Mon E	Tue D	Tue E	Wed D	Wed E	Thu D	Thu E	Fri D	Fri E	Sat D	Sat E	Sun D	Sun E	Total	Distance €
Base	1022	1359	1009	496	772	1325	731	586	980	1210	743	1125	877	0	12235	€ 12.235
I	1022	1357	1062	290	721	1328	720	562	826	1236	906	1015	833	249	12127	€ 12.127
II	1408	825	929	366	747	1379	911	570	829	1210	564	1179	833	0	11750	€ 11.750
III	834	1382	880	310	721	1315	708	754	829	1244	920	1011	810	0	11718	€ 11.718
IV	1054	1339	919	331	676	1203	740	648	865	1057	913	1073	884	84	11786	€ 11.786
V	1163	737	800	368	609	1269	716	577	798	1246	732	1058	820	335	11228	€ 11.228
VI	1139	935	834	258	585	1164	712	648	859	1198	808	1029	1082	0	11251	€ 11.251
VII	883	1227	880	331	567	1361	848	610	830	1246	775	1011	832	0	11401	€ 11.401
VIII	855	1170	941	242	540	1340	674	646	852	1258	704	995	838	0	11055	€ 11.055
IX	1230	798	823	322	572	1228	642	1104	697	1676	679	995	942	0	11708	€ 11.708
X	869	1226	909	244	682	1300	809	673	858	1283	827	950	832	0	11462	€ 11.462
XI	695	1321	725	433	562	1350	685	926	250	1877	633	1117	918	0	11492	€ 11.492
XII	919	1358	936	449	690	1371	717	586	870	1210	813	1032	892	126	11969	€ 11.969
XIII	1143	866	827	415	709	1287	668	627	839	1155	778	1010	976	0	11300	€ 11.300
XIV	1066	871	934	207	912	1152	440	883	691	1236	585	1078	1190	0	11245	€ 11.245
XV	979	1037	838	366	700	1343	717	586	764	1340	903	1008	863	0	11444	€ 11.444
XVI	909	1134	615	512	950	989	698	818	385	1360	555	1335	990	0	11250	€ 11.250
XVII	726	1283	737	405	862	906	647	875	279	1572	646	1098	835	0	10871	€ 10.871

Table 38: Average delivery size per shift

Experiment	Mon D	Mon E	Tue D	Tue E	Wed D	Wed E	Thu D	Thu E	Fri D	Fri E	Sat D	Sat E	Sun D	Sun E	Total
Base	7,6	8,9	7,5	9	7,7	9,8	8	9,1	8,4	10	8,7	10	8,4		8,7
I	10	10,8	5,3	3,6	5	8,5	6,3	6,9	9,5	10,2	9,2	10,8	12	5	8,6
II	7,3	8,1	7,1	9,9	8,8	11,1	8,3	11,7	6,3	9,3	8,7	9,2	10,9		8,8
III	7,1	9,1	6,9	12,1	7,8	10	7,9	10,4	6,6	9	8,6	9,8	11,3		8,8
IV	10,6	10,5	7,4	7,1	6	8,2	7,1	9,6	9,2	10,1	9	9,9	10,2	6	8,9
V	7,7	7,7	6,6	10,4	8,8	10,3	8,3	12,5	6,2	9,1	9,5	10,9	11,2	3,8	8,9
VI	6,9	10,1	7,5	9,8	7,9	8,5	8,5	13,7	5,3	9,7	9,3	11,9	10,6		9,1
VII	7,9	8,6	7,7	9,4	9,5	9,4	6,9	10,4	9	10,1	8,6	9,3	11,2		9
VIII	7,8	8,8	8,2	8,9	9,5	9,8	7,7	12,1	8,1	10	7,8	9,8	10,5		9,1
IX	7,8	9,4	7,8	9,6	9,1	9,4	8,7	9,8	6,2	9,5	7,7	8,9	7,5		8,6
X	8,6	8,9	8,5	8,3	8,3	9,3	6,9	10,9	5,8	9,8	8,8	9,4	11,4		8,8
XI	9,7	8,5	7,2	9,3	9,6	9,2	7,3	11,6	6	8,8	9	9,8	9,2		8,9
XII	6,5	7,2	7,2	9,2	8,8	9,3	8	11,5	5,8	9,1	8,3	10,9	10	8,7	8,5
XIII	7,6	7,9	7,1	8,8	8,9	9,8	8,8	11,1	6,1	9,4	8,7	11,4	11,6		9
XIV	7,9	8	7,7	8,3	10,6	9,5	8,4	10	6,4	8,9	9,2	10,8	11,1		9
XV	7	7,7	7	10,5	8,7	9,7	7,9	11,5	5,6	9	8,6	11,2	11,4		8,8
XVI	6,3	8,2	7	9,5	9,8	9,7	8,5	10,6	4,6	8,7	9,6	10,4	11,3		8,9
XVII	8,3	9	6,9	10,9	8,4	9,6	7,3	12,3	5,4	10	8,8	9,1	10,7		9,1

Table 39: Impact on fleet size and corresponding costs / Centre of Gravity

Experiment	Fleet impact	Truck €	CoG
Base	0	€0	7,10
I	0	€0	7,30
II	1	€2000	7,19
III	0	€0	7,16
IV	0	€0	7,04
V	0	€0	7,40
VI	1	€2000	7,48
VII	0	€0	7,25
VIII	0	€0	7,24
IX	0	€0	7,08
X	-1	- €2000	7,20
XI	-2	- €4000	7,29
XII	0	€0	7,36
XIII	0	€0	7,47
XIV	0	€0	7,52
XV	-1	- €2000	7,37
XVI	-2	- €4000	7,54
XVII	-2	- €4000	7,34

Results operational policy

Table 40: Routes per shift

Experiment	Mon D	Mon E	Tue D	Tue E	Wed D	Wed E	Thu D	Thu E	Fri D	Fri E	Sat D	Sat E	Sun D	Sun E	Total
Base	4	6	3	2	3	6	3	3	3	5	4	4	4	0	50
Base 22	4	6	3	2	3	6	3	2	3	5	4	4	4	0	49
Base 27	4	6	3	2	3	6	3	3	3	5	4	4	4	0	50
Base 30	4	6	3	2	4	6	3	3	3	5	4	4	4	0	51
XVI	3	5	2	3	4	5	3	4	1	5	3	6	5	0	49
XVI 22	3	6	2	2	3	6	3	4	1	6	3	4	4	0	47
XVI 27	3	6	2	2	3	6	3	4	1	6	4	4	4	0	48
XVI 30	3	6	4	0	5	4	3	4	2	6	4	4	5	0	50

Table 41: Delivery units per shift and warehouse related costs

Experiment	Mon D	Mon E	Tue D	Tue E	Wed D	Wed E	Thu D	Thu E	Fri D	Fri E	Sat D	Sat E	Sun D	Sun E	Total	Warehouse €
Base	183	284	164	117	162	293	175	137	151	279	192	229	201	0	2567	€ 11.860
Base 22	156	283	141	63	137	245	144	114	114	245	165	157	186	0	2150	€ 9.910
Base 27	169	307	147	69	157	269	164	134	142	267	205	234	217	0	2481	€ 11.502
Base 30	216	343	163	118	197	313	177	155	157	279	201	240	224	0	2783	€ 12.924
XVI	176	213	119	123	225	263	179	180	51	260	173	323	282	0	2567	€ 12.182
XVI 22	132	276	109	64	107	282	144	154	59	318	149	176	185	0	2155	€ 9.870
XVI 27	142	282	119	68	122	308	161	174	58	358	180	239	239	0	2450	€ 11.324
XVI 30	179	344	225	0	282	228	164	200	68	349	203	227	240	0	2709	€ 12.486

Table 42: Utilization per shift

Experiment	Mon D	Mon E	Tue D	Tue E	Wed D	Wed E	Thu D	Thu E	Fri D	Fri E	Sat D	Sat E	Sun D	Sun E	Total
Base	76%	79%	91%	98%	90%	81%	97%	76%	84%	93%	80%	95%	84%		85,6%
Base 22	65%	79%	78%	53%	76%	68%	80%	95%	63%	82%	69%	65%	78%		73,1%
Base 27	70%	85%	82%	58%	87%	75%	91%	74%	79%	89%	85%	98%	90%		82,7%
Base 30	90%	95%	91%	98%	82%	87%	98%	86%	87%	93%	84%	100%	93%		90,9%
XVI	98%	71%	99%	68%	94%	88%	99%	75%	85%	87%	96%	90%	94%		87,3%
XVI 22	73%	77%	91%	53%	59%	78%	80%	64%	98%	88%	83%	73%	77%		76,4%
XVI 27	79%	78%	99%	57%	68%	86%	89%	73%	97%	99%	75%	100%	100%		85,1%
XVI 30	99%	96%	94%		94%	95%	91%	83%	57%	97%	85%	95%	80%		90,3%

Table 43: Total duty time per shift and corresponding costs

Experiment	Mon D	Mon E	Tue D	Tue E	Wed D	Wed E	Thu D	Thu E	Fri D	Fri E	Sat D	Sat E	Sun D	Sun E	Total	Driver €
Base	35,3	45,5	32,1	16,7	27,7	43,5	28,0	22,9	30,2	40,3	30,4	35,2	31,7	0,0	419,517	€ 14.520
Base 22	33,6	45,7	30,6	14,0	25,7	43,5	26,8	18,5	28,8	40,6	29,6	33,8	30,1	0,0	401,167	€ 13.889
Base 27	33,6	45,8	30,9	14,2	26,9	44,0	27,6	21,7	27,6	40,5	30,7	35,5	31,8	0,0	410,583	€ 14.264
Base 30	35,9	46,4	30,9	16,3	34,8	44,4	28,0	22,7	30,1	40,5	30,6	35,7	31,9	0,0	428,2	€ 14.798
XVI	31,2	37,9	20,4	19,9	33,6	35,4	27,5	29,1	12,3	44,3	23,7	46,8	37,7	0,0	399,633	€ 14.176
XVI 22	22,7	46,8	23,4	12,2	20,1	45,5	25,8	29,0	10,0	56,5	25,0	34,7	31,0	0,0	382,6	€ 13.304
XVI 27	23,9	46,1	24,8	12,6	21,5	45,1	27,4	29,6	9,9	58,1	27,8	36,2	33,2	0,0	396,05	€ 13.837
XVI 30	25,3	51,1	42,0	0,0	42,0	31,5	28,2	30,5	14,4	55,6	32,9	35,5	36,3	0,0	425	€ 14.864

Table 44: Total traveled distance per shift and corresponding costs

Experiment	Mon D	Mon E	Tue D	Tue E	Wed D	Wed E	Thu D	Thu E	Fri D	Fri E	Sat D	Sat E	Sun D	Sun E	Total	Distance €
Base	1022	1359	1009	496	772	1325	731	586	980	1210	743	1125	877	0	12235	€ 12.235
Base 22	1000	1359	955	437	754	1325	728	503	980	1234	743	1125	805	0	11948	€ 11.948
Base 27	972	1321	955	437	768	1325	731	585	852	1235	743	1125	862	0	11911	€ 11.911
Base 30	1017	1359	932	492	1002	1316	725	586	980	1218	743	1125	862	0	12357	€ 12.357
XVI	909	1134	615	512	950	989	698	818	385	1360	555	1335	990	0	11250	€ 11.250
XVI 22	607	1355	728	301	562	1350	686	922	250	1819	636	1117	834	0	11167	€ 11.167
XVI 27	658	1319	793	303	556	1366	715	922	250	1839	657	1117	914	0	11409	€ 11.409
XVI 30	685	1434	1203	0	1077	898	776	902	365	1748	881	1119	936	0	12024	€ 12.024

Table 45: Average delivery size per shift

Experiment	Mon D	Mon E	Tue D	Tue E	Wed D	Wed E	Thu D	Thu E	Fri D	Fri E	Sat D	Sat E	Sun D	Sun E	Total
Base	7,6	8,9	7,5	9	7,7	9,8	8	9,1	8,4	10	8,7	10	8,4		8,7
Base 22	6,8	8,8	6,1	5,3	6,9	8,2	6,5	7,6	6,3	8,8	7,5	6,8	7,8		7,4
Base 27	7,3	9,6	6,4	5,8	7,5	9	7,5	8,9	7,9	9,5	9,3	10,2	9		8,5
Base 30	9	10,7	7,4	9,1	9	10,8	8	10,3	8,7	10	9,1	10,4	9,3		9,5
XVI	6,3	8,2	7	9,5	9,8	9,7	8,5	10,6	4,6	8,7	9,6	10,4	11,3		8,9
XVI 22	7,8	7,7	6,1	7,1	7,6	8,3	7,2	9,1	5,9	8	7,5	7,3	7,1		7,6
XVI 27	7,9	8,1	6,6	6,8	8,1	9,1	7,7	10,2	5,8	9	9	10	9,6		8,5
XVI 30	9,9	9,3	7,8		9,1	11,4	8,2	11,1	6,2	8,9	8,5	9,5	9,2		9,1

Table 46: Impact on fleet size and corresponding costs / Centre of Gravity

Experiment	Fleet impact	Truck €	CoG
Base	0	€0	7,10
Base 22	0	€0	7,01
Base 27	0	€0	7,24
Base 30	0	€0	6,97
XVI	-2	- €4000	7,54
XVI 22	-2	- €4000	7,24
XVI 27	-1	- €2000	7,50
XVI 30	0	€0	7,10