

How VMI Can Be Successful in Gas Distribution

A Solution Methodology for the Inventory Routing Problem in Gas Distribution

P.J.H. Hulshof, MSc
August, 2008, University of Twente



Committee

E.W. Hans, PhD, MSc	A. Hoendervoogt, MSc
A.L. Kok, MSc	J. Meesters, MSc



Awarded with the KVG Dutch Gas Industry
Master Thesis Prize 2010 (2nd Prize)

Preface

I have experienced my student time as a marvellous time, studying, travelling, starting a company, participating in student organisations, and meeting many different people and cultures. I thank all my friends and family for this wonderful experience.

Graduation is a rough experience to many people, but I experienced it as an interesting period full of exploration as well. This was mainly because ORTEC provided me with an interesting project and all the support that I needed. ORTEC made it possible that I could visit two planning sites in the Netherlands and in the United Kingdom to learn more about the practice of gas distribution. I thank ORTEC for offering me this research project, and I especially thank Arjan Hoendervoogt en Janneke Meesters for their support. Additionally, I thank the complete department of ORTEC Oil, Gas and Chemicals, who I bothered with many questions, and Joaquim Gromicho and Goos Kant from ORTEC for the meetings we had about the algorithm design.

Next to ORTEC, I was also motivated by the enthusiasm of Erwin Hans and Leendert Kok, my supervisors at the University of Twente. I thank them for their unbiased views and challenging propositions that always kept me thinking. Additionally, I thank Professor Dror from the University of Arizona for sending me his articles by mail.

Graduation is not always a sweet journey, and I thank my girlfriend, Kirsten, and my roommate, Jesse, for helping me through the tougher periods.

It is very sad that my grandfather passed away three months before my graduation, he would have been as proud as the rest of my family is now. He inspired, and always will inspire me to study more and to reach for higher objectives. I thank my father, Jacques, my mother, Ineke, my sister, Annelies, and the rest of my family for all their support during my studies, and for teaching me more than I will ever learn somewhere else.

Peter Hulshof
Amsterdam, August 2008
peter.hulshof@gmail.com

Abstract

Introduction

This research is done at the department Oil, Gas, and Chemicals at ORTEC, a planning software company. We study the distribution of gas to commercial and residential customers that are not connected to a network of gas pipelines. These customers receive gas deliveries under a Vendor Managed Inventory (VMI) contract, which gives gas companies the flexibility to determine when and what volume to deliver, and what routes to choose. The decision problem that is associated with VMI for a large set of customers is the Inventory Routing Problem (IRP). Additionally, gas companies want to control the effects of the large seasonal peak in gas demand, to use the available resources efficiently. This research assumes customer usage to be deterministic, and we develop a solution for a region with multiple depots and vehicles with varying capacity (heterogeneous fleet).

Objective

To design a solution methodology to minimize distribution costs in the IRP for gas distribution, and mitigate the seasonal peak in customer deliveries.

We propose a solution methodology that increases the *volume per kilometre*, since it is an important indicator of distribution costs. Additionally, we balance the delivery volume in the planning period, to use the available resources efficiently. To mitigate the seasonal peak, we balance the delivery volume over a relatively long period, so that the workload is more equally divided over the year.

Solution

An algorithm is developed to select a certain delivery day in the planning period for every customer. The algorithm focuses on finding delivery days for customers that can receive a relatively large delivery compared to the customer's capacity, while minimizing total travel distance. Customers that require a delivery in the planning period must be planned, and the customers that do not require a delivery in the planning period are planned according to the impact of the delivery on total travel distance and on future planning periods. The delivery volume is balanced according to the available vehicle capacity to efficiently use the available workforce and to smoothen the delivery volume over the course of a year.

Results

Actual delivery data from a large gas company are used to test the algorithm in a planning period of seven days. The computational experiments show that the solution increases the delivered *volume per kilometre* by more than 21%. The delivery volume is balanced on the short-term and long-term, and is responsive to changes in the vehicle capacity in the planning period.

Conclusions

The solution decreases the costs for gas distribution, and requires an acceptable computation time. The long-term balance in delivery volume flattens the customer delivery curve, and thus helps in mitigating the seasonal peak.

Contents

Preface.....	2
Abstract	3
Definitions	6
1 Introduction.....	7
1.1 Thesis structure	7
1.2 ORTEC	7
1.3 Problem description	8
1.4 Objective	9
1.5 Scope.....	10
1.6 Research questions.....	10
2 Context analysis.....	11
2.1 The Inventory Routing Problem	11
2.1.1 The IRP in an example.....	12
2.2 Gas distribution.....	13
2.2.1 The economics of gas distribution	14
2.2.2 Solution requirements for gas distribution	15
2.3 ORION and SHORTREC	15
2.4 The background solution: Period Scheduler.....	16
2.4.1 Background solution analysis	17
2.5 Solution requirements	17
3 Literature analysis	18
3.1 Key principles in solving the IRP.....	18
3.2 Solution methodology for the IRP	19
3.3 Performance measurement for the IRP.....	21
3.4 Seasonal peak problem.....	21
3.5 Decision models to optimize the IRP solution	22
3.5.1 Integer Linear Programming	22
3.5.2 Mixed Integer Linear Programming.....	23
3.5.3 Dynamic Programming	23
3.6 Main conclusions from the literature	23
4 Solution.....	24
4.1 Algorithm design.....	24
4.1.1 Solution requirements.....	24
4.1.2 Assumptions	24
4.2 Algorithm description	25
4.2.1 Algorithm overview	25
4.2.2 The algorithm in detail	26
4.2.2.1 Step 1: Customer selection	26
4.2.2.2 Step 2: Flexible clustering.....	27
4.2.2.3 Step 3: Generation of schedules	28
4.2.2.4 Step 4: Schedule assignment with an ILP for seed customers	29
4.2.2.5 Step 5: Schedule assignment with an ILP for clustered customers	31
4.3 Alternative algorithm designs.....	34
4.4 Distance types.....	35
4.5 Parameters of the algorithm	35

5	Computational results	36
5.1	Experimental design	36
5.1.1	Algorithm settings	36
5.1.2	Dataset	36
5.1.3	Assumptions	38
5.2	Comparison base scenario and algorithm	38
5.2.1	Weekly comparison	38
5.2.2	Workload balance	39
5.2.3	Different vehicle capacity set-up	41
5.3	Sensitivity analysis	42
5.3.1	Setting α in customer selection	42
5.3.2	Long-term workload balance	44
5.3.3	Short-term workload balance	45
5.3.4	Distance types	45
5.4	Testing alternative algorithm designs	46
6	Conclusions and recommendations	47
6.1	Conclusions	47
6.2	Recommendations	48
6.2.1	Long-term recommendations	48
6.2.2	Short-term recommendation	48
6.3	Scientific contributions	48
6.4	Future research	49
	References	50
	Appendices	52
	Appendix A: The Period Scheduler: Seed selection procedure	53
	Appendix B: The Period Scheduler: ILP model	54
	Appendix C: Extended literature summary	55
	C.1 Key principles in solving the IRP	55
	C.2 Solution methodology for the IRP	57
	C.3 Performance measurement for the IRP	63
	C.4 Seasonal peak problem	64
	Appendix D: Forecasted and actual customer usage curves	65
	Appendix E: Symbols used in the mathematical formulations	68
	Appendix F: Workload calculation: Step 4	70
	Appendix G: Workload calculation: Step 5	72
	Appendix H: Workload calculation: Improving the balance	75
	Appendix I: Alternative designs: Balance customers	76
	Appendix J: Alternative designs: Different objective functions	80
	Appendix K: Alternative designs: ILP with combination check	82
	Appendix L: Settings used in the computational tests	83
	Appendix M: Dataset overview	84
	Appendix N: Comparison actual routes and SHORTREC routes	85
	Appendix O: Computational results: The weekly comparison	86
	Appendix P: Computational results: The working time balance	87
	Appendix Q: Computational results: A different vehicle set-up	88
	Appendix R: Computational results: α parameter	89
	Appendix S: Computational results: Bandwidth parameter	90
	Appendix T: Computational results: Distance type parameter	91
	Appendix U: Computational results: Alternative designs	93

Definitions

This research is about propane gas, and we refer to propane gas simply with **gas**. Propane gas is used for heating residential and company buildings.

An **order** is an intention to buy or sell a certain amount of goods. An order can be **pulled**, meaning that an order is an instruction from the buyer that he wants to buy. If an order is **pushed**, an order is forecasted by the seller, and delivered to the buyer. A **delivery** is a planned and scheduled order, waiting to be delivered, and an **actual delivery** is a completed delivery.

Volatility is the standard deviation in demand. A **volatile good** refers to a good that has a relatively unstable demand curve with many peaks. The **safety stock** is a certain amount of stock that is kept to buffer against stock outs. In a highly volatile market, safety stock should be chosen to be relatively high, because peaks in demand are larger, and thus the risk of a stock out is larger (Brown et al., 2002). A **seasonal profile** reflects the volatility of a good in a year, for instance an increase in demand in the winter period.

Workload balancing has the objective to balance the volume of work over periods of time, to have a relatively constant flow of work, and to mitigate the influence of volatility.

A **client** is a customer of ORTEC. In this thesis the gas companies are clients. A **customer** needs to be replenished by ORTEC's client. A **location** is a customer that needs to be visited in the routing problem.

A **'must-go' customer** requires a delivery in the planning period, because the customer will reach safety stock otherwise. A **'may-go' customer** will not reach safety stock in the planning period, but may be delivered to create efficient routes.

The **visit frequency** is the number of visits per period that are needed to replenish a certain customer. **Slow-movers** are customers that have a relatively low visit frequency, for instance once per year. **Fast-movers** have a relatively high visit frequency, for instance once per week.

A **truck** is a lorry without a **trailer**. The trailer is a compartment to store and transport goods in. In our thesis, the trailer is a large tank, which is pulled by a truck. A **driver** is an employee that drives the truck to its destination. A **vehicle** is a combination of a truck, a trailer, and a driver. Vehicle capacity may vary, and a certain **vehicle type** indicates a vehicle with a certain load capacity.

A **depot** is a central location in a region with a large stock of gas, and where vehicles can load and unload. In this thesis, it is assumed that depots have infinite stock.

A **trip** is a set of orders that can be delivered by a vehicle in one haul, with the restrictions on the driver's work time and the vehicle's capacity taken into account. A vehicle can perform multiple trips per day.

A **network** is a set of locations that are called **nodes**, and paths between these nodes that are called **arcs**. An arc between two locations can for instance be depicting the distance between these two locations.

1 Introduction

Gas is a commodity, which is of crucial importance to our daily lives. Gas is used for heating our houses or work environment. To deliver gas to their customers, gas companies make a complex set of decisions in exploration, production, and distribution. Although most customers are connected to large national networks of pipelines, some customers are not. For these customers, gas is supplied by trucks and stored in large tanks at the customer's location. To make transportation by truck more efficient, gas companies supply gas to their customers according to a Vendor Managed Inventory (VMI) agreement. With VMI, the decision to deliver gas lies with the gas company, and not with the individual customer.

When a customer agrees with a VMI contract, a gas company can choose how often, when, and in what quantities the customer is replenished. In return, the gas company ensures that the customer will never run out of stock. In order to do this, a gas company needs to forecast and monitor the customer's gas usage. This gas usage is influenced by the outside temperature, meaning that usage rises when temperature is low and vice versa. Due to this dependence, there is a high peak in gas demand in the winter months. This increases pressure on the gas company's resources, and eventually leads to stock outs at customers in periods of high demand.

This research proposes a solution methodology to find a better distribution strategy when a gas company has the freedom that comes with VMI, and to mitigate the delivery peak in the winter months. The structure and lay-out of the thesis is discussed in Paragraph 1.1. This research has been done at ORTEC; we discuss the company in Paragraph 1.2. Paragraph 1.3 describes the problem, and Paragraph 1.4 states the research objective. The scope of the project is discussed in Paragraph 1.5, and Paragraph 1.6 clarifies the research questions.

1.1 Thesis structure

- Chapter 1 states an overview of the research. The company ORTEC is described, and we give a description of the problem.
- In Chapter 2, we analyze the context of the problem. We discuss the main problem, gas distribution, the involved software suites, and the background solution currently in place.
- Chapter 3 describes the literature that applies to the research questions in this thesis.
- Chapter 4 describes our solution, based on the information from Chapters 2 and 3.
- In Chapter 5, we discuss the computational results and a sensitivity analysis.
- In Chapter 6, we discuss the conclusions and the recommendations. We also discuss the scientific contributions of our report, and state ideas for future research.

1.2 ORTEC

This research is done by the Oil, Gas and Chemicals department of ORTEC. ORTEC is a consultancy company founded in 1981. The company is specialized in advanced planning and scheduling solutions, and is also active in financial risk management. ORTEC has 15 offices in Europe and North America and more than 800 customers worldwide. The fast-growing company employs over 700 people and delivers products in various industries: aviation, healthcare, transport, oil, gas, etc. The Oil, Gas and Chemicals department in the Dutch headquarters in Gouda, consists of 15 employees and is involved in large international projects at major oil and gas companies.

1.3 Problem description

Gas companies supply gas to their customers according to a VMI agreement. The gas company can optimize its decisions on the condition that at least a minimum stock of gas is always available to the customer. This minimum stock is called the safety stock, and it serves as a buffer to prevent stock outs. Because of VMI, gas companies forecast the customer usage, and have more freedom in creating their distribution plan. Additionally, customers do not need to dedicate resources to inventory management.

ORTEC provides a software solution called ORION to support gas companies in forecasting gas usage with historical delivery data. From the forecasted usage, orders are generated which are exported to SHORTREC, a software product that creates solutions for large instances of the Vehicle Routing Problem (VRP). The VRP is the problem of selecting a set of routes for a certain vehicle fleet in such a way that all orders are delivered at minimal costs (Dantzig and Ramser, 1959).




Figure 1.1. The two distinct phases in planning and scheduling of all customer deliveries.

ORION and SHORTREC decompose the problem of planning and scheduling into two clear phases that are illustrated in Figure 1.1.

ORION plans the delivery date of an order with an algorithm that minimizes the number of visits. This means a customer is served as late as possible, just before the customer reaches safety stock. According to Campbell and Savelsbergh (2004), this approach maximizes the volume deliverable to a customer, but may not create an optimal solution on the long-term, since it does not recognize geographical synergies between customers. The savings in synergy may be higher than the savings in minimizing the number of visits, for instance if two customers are located not far from each other, but should be visited on a different day according to ORION. Figure 1.2 illustrates this problem.

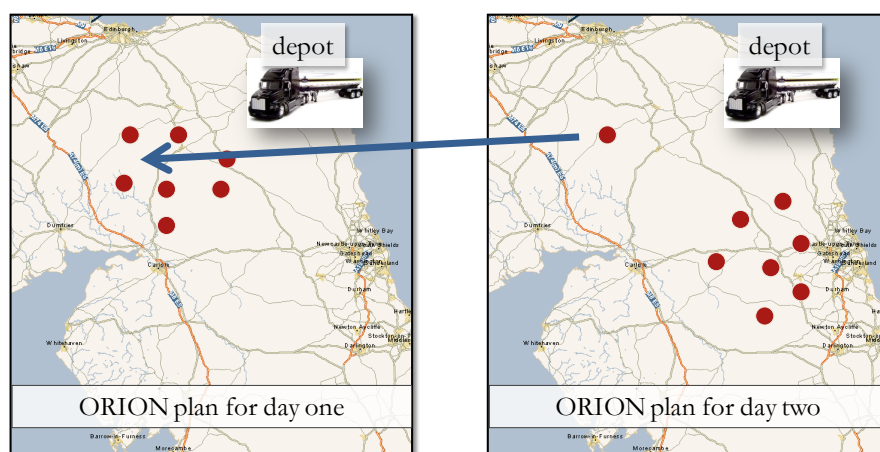


Figure 1.2. The blue arrow indicates that the order (in red) is more efficiently delivered, when it is delivered on the same day as customers in close vicinity.




Figure 1.3. The actual weekly delivery volumes of gas in the Northern-England region (NE-region) clearly show the seasonal peak in demand in the winter months

Figure 1.3 illustrates a large and long lasting peak in the demand for gas in the winter months; all other peaks and disturbances are explained in Paragraph 2.2. On average, the volume delivered to customers in the winter months is 1,5 times as high as the amount delivered in the summer months. The workload balance over a full year is offset by the large peak in demand in the winter months. This leads to stock outs at customers during the winter months and, in combination with a fixed fleet over the year, to unused trucks in the summer months. If we want to balance the workload over a full year, it is important to consider the expected demand in a longer period in the short-term plan. Currently, there is no available method in ORION or SHORTREC to plan and schedule orders, taking future total demand in the region into account.

Additionally, we study the problem of unstable workload balance on the short-term. To explain this problem, we provide an example: if fewer orders are planned on Wednesday than on Friday, we do not need all trucks on Wednesday. Then on Saturday, we need an emergency truck, since not all orders can be delivered on Friday, and this results in high costs.

In the literature, the problem described in this paragraph is referred to as an Inventory Routing Problem (IRP). The IRP is similar to a VRP with two additional questions: (1) when to serve a customer, and (2) how much to deliver to this customer. Paragraph 2.1 describes the IRP in detail. We extend the IRP with the problem of the seasonal peak in demand.

1.4 Objective

We identify two separate problems, the IRP, and the problem of the seasonal peak in demand in the winter period. We state a research objective that is directly applicable to these problems:

To design a solution methodology to minimize distribution costs in the IRP for gas distribution, and to mitigate the seasonal peak in customer deliveries.

1.5 Scope

We examine the first of the two phases in Figure 1.1: the generation of orders. More specifically, we investigate the interaction in the first phase between the ORION and SHORTREC suites. We do not focus on the usage forecasting, but specifically on the translation of this usage forecast to orders. Additionally, our objective is not to find improvements in the optimization algorithms in SHORTREC. We exclude the problem of multiple tank compartments, which is common in the oil industry, and we investigate the situation with one compartment, because this setting is common in gas distribution.

1.6 Research questions

We set apart the seasonal peak problem from the IRP, because the seasonal peak problem is an extension of the IRP. Literature on the IRP with methods to solve the seasonal peak is not available.

The IRP

The answer to RQ1 analyzes the context of the IRP. RQ2 studies the available methods in theory and practice to solve the IRP. RQ3 connects the IRP to a short-term workload balance. RQ4 relates to the implementation and evaluation of our solution.

- RQ1. How can we describe the classical IRP? [Paragraph 2.1]
 - a. What restrictions follow from the gas market? [Paragraph 2.2]
 - b. What is the background solution? [Paragraph 2.4]
 - c. What are the solution requirements? [Paragraph 2.5]
- RQ2. How can we solve the classical IRP? [Paragraph 3.2]
 - a. When to serve a customer? [Paragraph 3.1]
 - b. How to reflect the long-term effect of a short-term decision? [Paragraph 3.1]
 - c. How can we apply this method to our extension of the IRP? [Chapter 4]
 - d. How to include the multi-depot problem? [Chapter 4]
 - e. What decision model can we use? [Paragraph 3.5]
- RQ3. How to balance the short-term workload in the IRP? [Paragraph 3.2]
 - a. What criterion for balancing the workload should we use? [Paragraph 3.2]
- RQ4. How can we apply our solution to ORION and SHORTREC? [Paragraph 6.2]
 - a. What is the function of ORION and SHORTREC? [Paragraph 2.3]
 - b. How should the performance of our solution be measured? [Paragraph 3.3]
 - c. What are the best settings for our solution? [Paragraph 5.3]

The seasonal peak

The answer to RQ5 results in a solution to the seasonal peak problem, since the workload is better balanced over a full year.

- RQ5. How can we balance the workload in the IRP for a full year? [Paragraph 3.4]

2 Context analysis

This chapter discusses the background of the problem. Paragraph 2.1 describes the IRP as stated in the literature. Paragraph 2.2 discusses the gas market and its properties. Paragraph 2.3 goes into ORION and SHORTREC, and Paragraph 2.4 describes the current solution for the problem: the Period Scheduler. Paragraph 2.5 gives a set of solution requirements to guide our literature research in Chapter 3.

2.1 The Inventory Routing Problem

Campbell and Savelsbergh (2004) describe the IRP as the distribution of a single product from a single depot, to a set of N customers over a given planning horizon of length T , possibly infinity. Customer i consumes the product at a given rate U_i (volume per day), and its storage capacity is given by C_i . The start inventory at customer i is I_i^0 at day t is 0. A fleet of M homogeneous vehicles, meaning the vehicles are all the same, is available for the distribution of the product. The vehicles have a fixed capacity Q . The objective is to minimize the distribution costs without causing stock outs at the customers, which is the same as the objective of VMI. In the IRP, three decisions play an important role:

- (1) When to serve a customer?
- (2) How much to deliver to a customer?
- (3) What delivery route to use?

When the size and date of the delivery are determined, the question what delivery route to use remains. This means a VRP for every day in the planning needs to be solved. Since the VRP is a generalization of the Travelling Salesman Problem¹ (TSP) and the TSP is NP-hard, the VRP is NP-hard as well (Garey and Johnson, 1979). An NP-hard problem is usually not solvable within polynomial time, meaning it can only be solved within reasonable time with a heuristic algorithm. Because of its complexity, the IRP is NP-hard as well, and we can only solve the IRP within reasonable time with a heuristic algorithm.

The classical IRP, which we described above, differs from our IRP on three points:

- (1) In our IRP the vehicle fleet is not homogeneous, i.e. we have trucks with varying capacity.
- (2) Some customers can only be visited by a certain type of vehicle in our IRP.
- (3) Multiple depots need to be considered when solving the problem in our IRP.

The IRP is a long-term dynamical control problem. It is dynamical, since a decision today affects the situation tomorrow. This long-term dynamical control problem is hard to formulate and to solve. Therefore, approaches to solve the IRP are focused on solving a short-term planning problem. Two questions are important in these approaches:

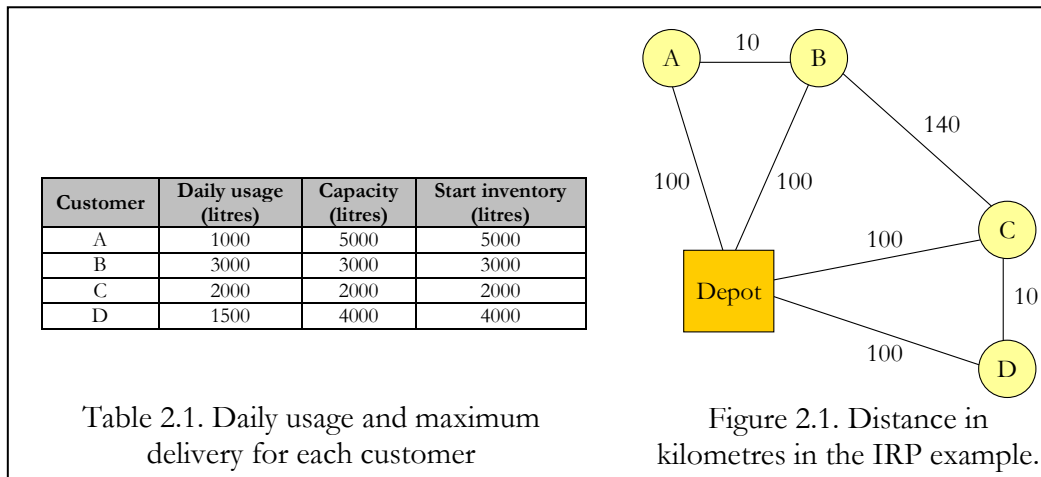
- (1) How to model and account for the long-term effect of short-term decisions?
- (2) What customers to include in the short-term planning period?

Following a short-term approach results in postponing as many deliveries as possible to a following planning period. This leads to problems in following planning periods.

¹ The objective in the TSP is to find an optimal route with minimal travel distance for a salesman, who must visit a number of locations and has to return to his starting point.

2.1.1 The IRP in an example

To illustrate the dynamics of the IRP, Bell et al. (1983) describe a simple example with four customers and one depot. Each customer has a fixed demand per day, a maximum capacity, and a starting inventory which is equal to its capacity; Table 2.1 displays the properties of all customers. Figure 2.1 illustrates the distances between the customers and the depot. There is one vehicle with a capacity of 5000 litres. Customers use the gas up to noon, and the vehicle starts delivering after noon. The vehicle can perform two trips per day, and we need a plan for the next two days.



The obvious solution to this problem would be to deliver Customers A and B on the first trip and Customers C and D on a second trip, both trips on both days. We would deliver 7500 litres per day, and drive 420 kilometres per day.

A better solution is to deliver to Customers B and C on the first day, and deliver to Customers A and B, and to Customers C and D on the second day. Figure 2.2 illustrates this, where a circle illustrates a trip and the different colours display different days. We can deliver volume for two days to Customers A and D and for one day to Customers B and C in this solution, in a total of three trips spread out over two days. We get an average of 380 kilometres and 7500 litres per day.




Figure 2.2. Geographical display of the better solution.

Customer	A	B	C	D	Total distance (km)	Total volume (litres)
Trip 1 (Day 1)	-	3000	2000	-	340	5000
Trip 2 & 3 (Day 2)	2000	3000	2000	3000	420	10000
Average per day					380	7500

Table 2.2. The better solution for the example.

2.2 Gas distribution

Gas is found in oil or gas fields, and it follows an intensive production process at a refinery before it is stored. From this storage, it can be supplied to depots in the market area by means of pipeline, boat, train or truck, in compressed or liquefied form. For customers that are not connected to the national grid, depots deliver gas by tank truck. This is also done for commercial customers, such as gas stations, who are selling gas to car owners. The customer stores the gas in a tank, and each customer requires a delivery between several times per week and once per year. Planners at the depots are responsible for the planning of these deliveries to make sure all customers have sufficient stock.




Figure 2.3. Forecasted weekly usage and actual weekly delivery volumes of gas in the Northern-England region in the weeks 2006-49 to 2007-48. The graph is different from Figure 1.3, because these data only represent customers that are in the VMI program.

Figure 2.3 illustrates the annual demand curve which shows a large seasonal peak in demand in the winter months. This is caused by additional heating due to the lower temperature. The graph is constructed from actual delivery data of three depots in the Northern-England region in the UK (NE-region). These three depots serve 3.188 customers, of which 60% are in the VMI program, and deliveries are controlled by the gas company. The other 40% of the customers are not in the VMI program, and the demand of these customers is not forecasted, meaning these customers order gas themselves. We use the data of the NE-region frequently in this research.

In the graph of Figure 2.3, all customers that have no forecasted usage are excluded. This concerns (1) new customers, since there are insufficient data to determine their usage rates, and (2) customers that are not in the VMI program. Grain farms are customers that are not in the VMI program, because their demand depends on many factors that are difficult to forecast. Grain farms have a relatively large demand in the months August and September, to dry the harvested grain.

The large drop in actual deliveries in week 51 is because of Christmas, which is an official holiday, and thus holds a lower amount of available resources. The large peak in actual deliveries in week 26 is due to a data error in ORION. Based on this data error, ORION underestimated the expected demand of the customers, and thus proposed expected delivery dates that were too late. To make up for this error, many deliveries had to be made in the weeks 25, 26, and 27. The forecasted usage in Figure 2.3 is determined in 2008, so it is based on corrected data. All other peaks and disturbances in the actual delivery curve can be explained by the volatility in planning the orders, or the volatility in available resources.

Due to safety regulations, the maximum capacity at a customer is 85% of the customer's tank capacity. An emergency delivery is scheduled, when a customer is out of stock, or nearly out of stock. The exact definition of an emergency delivery, is a delivery where the customer tank level is below 15%, and the order is added *after* a trip was created. There were 132 emergency deliveries in the period December 2006 until November 2007, and these are 0,8% of all orders in this period (15.929 orders). August and September have increased emergency deliveries, since grain farms have high and volatile demand in this period. Figure 2.4 illustrates the distribution of emergency orders over the year.




Figure 2.4. The number of emergency orders in the NE-region in the period December 2006 - November 2007.

Trips and restrictions

The average number of customers per trip in the NE-region is nine. This high number of stops per trip increases the complexity of calculating an optimal trip in the corresponding VRP. Many standard constraints apply to the calculation of these trips. Customers can be visited in time windows that are only bounded by the hours of daylight, and there is no need for customers to be available for the delivery. Additionally, two important restrictions apply in gas distribution: vehicle restrictions and equipment restrictions. Vehicle restrictions restrict the type of vehicle by which a customer can be visited. For instance, a customer that lives in an area that is very difficult to reach can only be visited by the smallest vehicle, since the other vehicles are too large for the roads to the customer. Equipment restrictions require certain equipment to be available in the vehicle to deliver to a customer. Although this restriction is not used in the NE-region, we add it to the solution requirements, since it is common in gas distribution.

2.2.1 The economics of gas distribution

The costs of gas distribution are mainly driven by volatility, travelled kilometres, and the number of visits to customers.

Volatility costs

Vehicle contracts are settled for a full year, so resources are fixed over the year. Volatility is a strong driver of cost, since the resources are fixed, but demand is not. Volatility is a cause for idle, or for emergency resources, and minimizing the effects of volatility is one of the objectives.

Vehicle fleet and travelled kilometres

In gas distribution, costs are specifically driven by the vehicle fleet, and the number of kilometres this fleet has to travel to deliver the gas.

Number of visits

The number of visits to a customer should be kept to a minimum, since a visit means additional kilometres, and thus additional costs. The optimal number of visits to a customer can be calculated by dividing the customer's usage by its capacity. This is optimal when a vehicle can only visit one customer per trip, but in gas distribution there are nine stops per trip. A customer can not be observed as an individual, but should be observed as a system, where customers in the vicinity influence the number of visits to this customer as well. This trade-off between minimizing the number of visits and minimizing the kilometres travelled is the crucial element of the IRP in gas distribution.

2.2.2 Solution requirements for gas distribution

The solution requirements for gas distribution are based on this paragraph:

- (1) Travel to different areas in the region on a single day to be able to respond to emergency orders.
- (2) Multi-depot approach, since a planner is usually responsible for a region with multiple depots in gas distribution.
- (3) Heterogeneous fleet approach, because there are different types and sizes of vehicles in real-life instances of the IRP in gas distribution.
- (4) Take in consideration vehicle and equipment restrictions.
- (5) Balance the workload to decrease volatility in deliveries.

2.3 ORION and SHORTREC

ORION and SHORTREC are two decision support software suites that communicate with each other by exchanging files. The suites are used by end-users at ORTEC's clients, and they find good results with relatively low computation times.

Planners plan gas deliveries for a certain planning period, usually one day. Based on historical data of tank measurements (dip levels) and deliveries (drop sizes), ORION forecasts gas usage in the future. With the information on usage, ORION generates the orders just before a customer reaches safety stock, with the objective to minimize the number of visits. Figure 2.5 illustrates the ideal curve of a customer's stock level.




Figure 2.5. The ideal curve of a customer's stock level to minimize the number of visits.

The vertical lines are deliveries, and the diagonal lines illustrate usage. The top and bottom horizontal lines illustrate the maximum stock level and the safety stock level.

ORION generates two types of orders; a 'must-go' and a 'may-go' order. A 'must-go' customer requires a delivery before a certain day in the planning period, because the customer will reach safety stock otherwise. A 'may-go' customer will not reach safety stock in the planning period, but may be delivered to create efficient routes.

Based on the orders from ORION, SHORTREC solves the VRPs for all orders that are planned on a specific day. SHORTREC produces a routing schedule that is created with a (1) constructive heuristic, and is optimized with (2) an improvement heuristic. A SHORTREC-user can select the savings-based or insertion heuristics as described in Poot et al. (2002). At the start of the algorithm, a trip is created for the customer that is farthest away from a depot. The main idea is to insert customers into this trip. These insertions come with certain costs in distance, and the customer with the lowest insertion costs is chosen. When a vehicle is filled, a new trip is started with the customer that is the farthest away from the depot. With SHORTREC, many restrictions can be taken into account: time windows, vehicle capacity, driver schedules, restricted routes, etc.

2.4 The background solution: Period Scheduler

The Period Scheduler is an algorithm that is created for the Period Vehicle Routing Problem (PVRP) for a large beverage company. This paragraph describes the Period Scheduler adjusted for gas distribution.

From the complete geographical area that is serviced by one depot, n separate seed points are chosen for every day of the planning period, so when there are five days, there are five seed points. These seed points are the centres of gravity of all customers, so they are fixed. The days of the planning period are assigned to the seed points according to their creation order, i.e. Seed 1 is assigned to Monday and Seed 2 to Tuesday, etc. See Appendix A for the detailed seed creation process.

Every customer is connected to a set of delivery scenarios, which contain percentages of the customer's tank size that can be delivered on a certain day in the planning period. The Period Scheduler assigns one of the scenarios to the customer by solving an Integer Linear Programming problem (ILP), which is given in Appendix B. An ILP is a Linear Programming problem (LP) with a restriction to the possible values of one or more variables. The restricted variables can only be integer values. An ILP is NP-hard, where an LP is not (Garey and Johnson, 1979). An ILP where the integer restriction is ignored is called an LP relaxation. The ILP in the Period Scheduler is solved by rounding the outcome of the LP relaxation to integers.

The ILP has the objective to minimize the driving distance between the customer and the seed points, which are connected to a certain delivery day. For a 'may-go' customer a comparison between a delivery this week and a delivery next week is incorporated in the cost function. The cost function for not visiting the 'may-go' customer this week, and forwarding the delivery into the future, is given in Equation 2.1.

$$c_t = \frac{A^3}{10^4 \cdot \sqrt[4]{B}} \quad \text{for } t = 6 \quad (2.1)$$

A is the percentage of customer's capacity that can be delivered the next week, if there is no delivery this week. B is the distance to the nearest 'must-go' customer in the current week. A low distance between the 'may-go' customer and its nearest 'must-go' customer increases the costs for forwarding the customer. Additionally, a relatively large delivery next week decreases the costs for forwarding the customer, since it has a relatively higher urgency.

The workload is balanced by delivery volume per day in the ILP. The average expected usage per week is calculated for 13 weeks into the future, and this is set as the workload for the current week. The end-user can set the percentage of the workload that has to be delivered on every day of the week. For instance, the end-user can plan fewer orders on Friday to reserve capacity for expected pull orders. If the expected demand has an increasing curve in the next 13 weeks, we will deliver more in the current week than only the expected demand for the current week, and thus we get a moving average on the long term that takes into account the expected demand for the next 13 weeks. Figure 2.7 illustrates the customer assignment to five seed points.

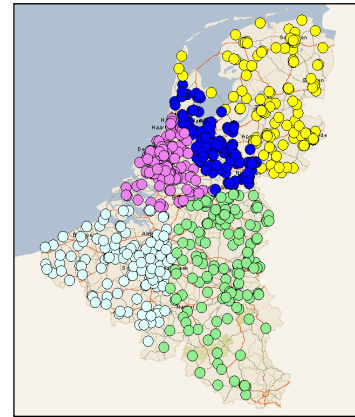


Figure 2.7. Colours mark the different days.

2.4.1 Background solution analysis

In our solution, we consider the following aspects:

- (1) The distance minimization of the Period Scheduler is a short-term perspective on cost. A long-term perspective on costs is only in the cost function for forwarding a ‘may-go’ customer by including the percentage of a ‘may-go’ customer’s capacity that can be delivered next week. Can this long-term approach also be incorporated in the cost function for delivery days within the planning period?
- (2) The Period Scheduler lacks flexibility, since only one seed point is selected for each day. In case of emergencies outside the seed point’s region, it is difficult to reschedule a truck.
- (3) The Period Scheduler can not cope with multiple depots in the planning region.
- (4) The Period Scheduler does not balance the workload considering vehicle capacity at a depot on a certain day in the planning period, but uses manual balancing parameters. It is beneficial to connect the availability of resources to the workload that should be planned on a specific day.

2.5 Solution requirements

We describe the main requirements for our solution, based on Chapters 1 and 2:

- (1) Relatively low computation times for real-life instances.
- (2) The workload is balanced to cope with volatility on the short-term and the long-term.
- (3) Travel to different areas in the region on a single day to be able to respond to emergency orders.
- (4) Multi-depot approach, since a planner is responsible for a region with multiple depots in gas distribution.
- (5) Heterogeneous fleet approach, because there are different types and sizes of vehicles in real-life instances of the IRP in gas distribution.
- (6) The solution is implemented and tested in the ORION and SHORTREC suites.

3 Literature analysis

This chapter analyzes the literature on the IRP and seasonal peak mitigation. For the reader that is merely interested in the main insights, a summary is given in Paragraph 3.6. Paragraph 3.1 discusses key principles, which form the basis for finding a solution methodology for any type of IRP. Additionally, we discuss five solution methodologies in Paragraph 3.2. Paragraph 3.3 discusses performance measurement for the IRP, and Paragraph 3.4 illustrates the seasonal peak problem in the literature. Paragraph 3.5 discusses the available decision models that are frequently used in the literature on the IRP. The reader is referred to Appendix C for the summaries of the articles discussed in this chapter.

3.1 Key principles in solving the IRP

The key principles form a better understanding of the trade-off between long-term and short-term costs. When deliver an order earlier than its optimal delivery point?

Dror and Ball (1987)

Dror and Ball (1987) propose a short-term planning solution, in which the long-term effect of a short-term decision is reflected. They find that the *relative delivery size* is a key indicator of the impact of a short-term decision on the long-term cost. If the relative delivery size is small, it means the delivery size is far from optimal, and costs are higher. Their idea for relative delivery size is also used in the Period Scheduler. The relative delivery is calculated by dividing the delivery for the customer by its capacity:

$$\frac{\text{Possible Delivery Volume}}{\text{Capacity}} \quad (3.1)$$

Campbell, Clarke, Kleywegt, and Savelsbergh (1997)

Campbell et al. (1997) propose two assumptions to guide decisions in solving the IRP:

- (1) Always try to maximize the quantity delivered, since this minimizes the number of visits to a customer on the long-term.
- (2) Always try to send out a full truck load, since this maximizes utilization.

Yugang, Haozun, and Feng (2008)

Yugang et al. (2008) approach transportation cost in a way that includes distance as *detour distance*, the additional distance that should be travelled if a customer is added to a trip. This is illustrated in Figure 3.1, a trip between the Depot and Customer 1 will be: $2 * 100 = 200$, and a trip including Customer 2 will be: $100 + 30 + 90 = 220$. This means we increase the distance of the trip with 20.




Figure 3.1. Detour distance is the extra distance by adding Customer 2 to the trip.

We state three measures for distance:

- (1) Euclidean distance, distance that is given by measuring a straight line on a map between two locations.
- (2) Real distance, distance that is given by measuring the minimum travel distance through a network of roads between two locations.
- (3) Detour distance, distance that is given by the additional travel distance when a certain customer is inserted in an existing trip. This is measured as a real distance.

3.2 Solution methodology for the IRP

Bell et al. (1983), Federgruen and Zipkin (1984), Golden et al. (1984), and Blumenfeld et al. (1987) are among the first authors to describe a solution methodology for the IRP. Since then, researchers have defined approaches to solve the IRP for different problem instances. The main differences in these approaches are:

- (1) A deterministic versus a stochastic approach, where customer usage is assumed to be either deterministic or stochastic. ORION uses deterministic customer usage.
- (2) A decomposition versus an integrated approach, where a decomposition approach tries to find a solution for the IRP in a phased approach. The first phase finds a timing and quantity for a delivery, and the second phase solves the resulting VRPs. The integrated approach deals with the decision when and how much to deliver, and the VRPs, at the same time.
- (3) Different decision models are used in the literature. Where most authors use an ILP to minimize cost or maximize revenue; other authors solve the IRP with Dynamic Programming (DP), or with heuristics.

All authors design solution methodologies for a homogeneous fleet. Although some methodologies cope with satellite reload facilities, no article has a multi-depot approach, where vehicles start and finish a day at multiple locations.

Golden, Assad, and Dahl (1984)

Golden, Assad, and Dahl (1984) state that the IRP is optimized along (1) a spatial dimension (distance) and (2) a temporal dimension (delivery timing). The authors use relative delivery size as a measure to reflect the temporal dimension. The distance is minimized and the relative delivery size is maximized. The authors use Equation 3.2 to select customers that have an opportunity to be delivered in a planning period.

$$\frac{\text{Possible Delivery Volume}}{\text{Capacity}} \geq \alpha \quad (3.2)$$

Dror and Trudeau (1988)

Dror and Trudeau (1988) investigate a stochastic IRP. The stochastic approach gives the opportunity to model route failures and stock outs in a simulation model. A route failure is a mismatch between the expected delivery volumes in a trip and the actual volumes, resulting in a trip with too little or too many customers. In our deterministic approach, it is difficult to test a solution for stock outs, since one can not simulate actual usage.

We deal with route failures by adding customers at the end of a trip with a flexible delivery. If the customers are not visited, there is no stock out at the customer, and if we do visit the customer, the delivery should be above a minimum economic delivery size. This method is similar to the method used in practice by the planners.

Bard, Huang, Jaillet, and Dror (1998)

Bard et al. (1998) discuss a decomposed approach for the IRP with satellite reload facilities. The bi-criteria approach they propose is applicable to our IRP, since it combines the maximization of the relative delivery size with the objective to minimize the distance travelled. Bi-criteria problems can be approached in two ways:

- (1) Both criteria are optimized simultaneously, by finding correct weights.
- (2) Optimize the criteria sequentially by first optimizing one criterion and then the other.

The authors use the second, separate optimization approach, and this approach is similar to the current set-up of ORION and SHORTREC, where ORION always generates an order when safety stock is almost reached, and SHORTREC minimizes distance. Since we want to have a combination of both, we need to minimize distance, and maximize the relative delivery in one evaluation step. The authors use volume as a factor to balance the workload over the week, and every day should have an equal volume to be delivered.

Campbell and Savelsbergh (2004)

Campbell and Savelsbergh (2004) propose a methodology that is appropriate for our IRP. The methodology consists of two phases that resemble the current planning and scheduling in ORION and SHORTREC.

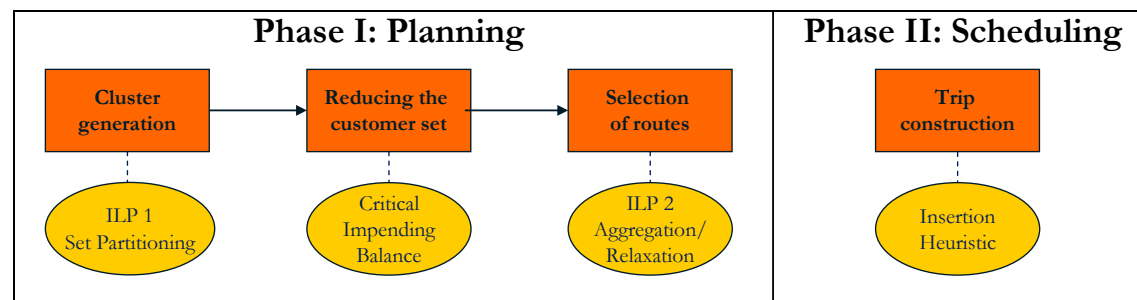


Figure 3.2. The solution methodology of Campbell and Savelsbergh (2004)

The authors reduce the problem size by creating clusters based on the knowledge that it should be possible to serve a cluster with one vehicle for a long period. They calculate the clusters once and re-calculate them when a new customer is added, or when customer usage patterns change. This form of ‘fixed pre-clustering’ requires large computation times, and is interesting in problem instances that are relatively small, since the customer set in these instances does not change often. With the high variability in customer usage curves and the many changes in the customer set in our IRP, we would update the clusters daily. We cluster in a different way, since we do want to use the distance minimization provided by clustering. ‘Must-go’ customers are the starting point for creating our clusters, since they form the basis for a solution methodology (Campbell and Savelsbergh, 2004). Clustering is also used by Jung and Mathur (2007) to minimize distances between customers before solving an actual VRP.

The authors reduce the customer set, which minimizes long-term costs and makes the algorithm faster. They only consider customers that can receive a certain relative delivery (Golden et al., 1984; Dror and Ball, 1987; Bard et al., 1998), or are very far from other customers. These customers are called impending and critical customers respectively. Additionally, b nearest neighbours to each of these customers are added as balance customers, to balance the workload equally over the days. The authors estimate and use working times to balance the workload in Phase I.

Kleywegt, Nori, and Savelsbergh (2004)

Kleywegt et al. (2004) use an integrated approach. The authors apply Dynamic Programming (DP), where the solution space is enumerated in an intelligent way, thereby creating good results. Although their results look promising, it is not useful yet. Kleywegt et al. (2004) apply their methodology to a problem where trips consist of three stops or less, and solution times already take multiple days. Trips in gas distribution have on average nine stops, and since their computation times grow exponentially with the number of stops per trip; we can not use their solution methodology. This was noted by Campbell and Savelsbergh (2004) as well.

3.3 Performance measurement for the IRP

The IRP is an NP-hard problem, and because of its complexity, we have no optimal solution to compare our heuristic with. Usually, a lower bound can be calculated for an NP-hard problem, which is a good solution to a simple derivation of the actual problem. A lower bound for our IRP will be very weak, because the large number of customers and the large number of stops per trip result in a prohibitively large number of options to calculate (Song and Savelsbergh, 2007). Song and Savelsbergh (2007) are the only authors that study performance measurement for the IRP, and they state that *volume per kilometre* is the best measure to compare IRP solution methodologies for equal problem instances.

Solution performance measures

Other authors in the IRP literature define transportation cost by the factors in Table 3.1.

Factor	Measure(s)	Available after activity
<u>Truck utilization</u>	Transported volume / Total used truck capacity	Scheduling
	Number of trucks needed	Scheduling
<u>Distance</u>	Kilometres travelled	Scheduling
<u>Volume</u>	Total planned volume	Planning
	Transport cost per volume	Scheduling
	Volume per order	Planning
	Volume per kilometre	Scheduling
<u>Driver cost</u>	Working time	Scheduling
	Driving time	Scheduling
<u>Number of customer visits</u>	Relative delivery (Delivery volume / Customer Capacity)	Planning

Table 3.1. Factors that influence transportation cost and measures to quantify them.

Some measures are available after the planning phase, and others require the scheduling phase to be completed as well.

3.4 Seasonal peak problem

Although some authors describe the seasonal demand peak (Dror and Ball, 1987), no author has ever written about a strategy to cope with this peak in gas distribution. Welch et al. (1971) propose five solutions for the seasonal peak in the production of gas, of which one is applicable to the distribution of gas as well.

Welch et al. (1971) state that a gas company should hold sufficient stock in peak periods, to cope with the higher demand. Thus if a gas company ensures that customers have sufficient stock during the peak period, it has to make fewer visits in the peak period. In the current situation, an order is generated as late as possible, and there is no connection to the peak in future demand of all customers.

We propose two methods to mitigate the seasonal peak in demand:

- (1) Take into account the future demand of all customers when assigning the orders to days. Use a workload constraint that balances the workload over a longer period than the planning period. This is also used in the Period Scheduler.
- (2) Ensure customers with a large flexibility do not require a delivery in peak times, by forcing orders for these customers in quieter times. For example, customers that need one delivery per year receive a delivery in summer, to avoid deliveries when the peak is at its highest, in the middle of winter.

We study the use of Option (1) to balance the workload over a full year, because this solution methodology is connected to the IRP. The graphs in Appendix D illustrate that (2) is not used, it is recommended for future research.

3.5 Decision models to optimize the IRP solution

The IRP is frequently solved with an ILP in the literature. Next to ILP, we discuss other models that are used: Mixed Integer Linear Programming (MILP), and DP.

3.5.1 Integer Linear Programming

An ILP is NP-hard, but can be solved fast with efficient heuristics. These heuristics are found in programs like CPLEX and AIMMS, but these programs can not be implemented in our solution, since their license costs are relatively high. We must use a solver that has no license costs. To find alternative methods, we discuss branch-and-bound, Lagrangian relaxation, and rounding.

Branch-and-bound

In branch-and-bound (Winston, 1994, *page 502-532*) the feasible solutions to an ILP are systematically enumerated, such that the optimal integer solution is found. This is done by solving an LP relaxation, and setting the solution to the upper bound of the problem (in a maximization problem). By making use of a tree-like setup, and determining the upper bound of each node in the tree, one can determine if it is still useful to evaluate the branches that lie below a node. We can only use branch-and-bound for small problem instances, since our solver uses prohibitively long computation times for larger instances.

Lagrangian relaxation

The main idea behind Lagrangian relaxation of an ILP is to incorporate the interfering constraints into the objective function, with a certain weight or penalty if the constraint is not satisfied. The problem can be solved by finding the best weights and the lowest upper bound for the maximization problem. We do not use this method in our algorithm, since the application of Lagrangian relaxation is studied in a future graduation thesis at ORTEC.

Rounding

A very fast approach to solve an ILP is to round the solution of the connected LP relaxation (Winston, 1994, *page 466*). In an LP relaxation, the integer constraint is dropped from an ILP, which makes it much easier to find a solution. Rounding the LP relaxation for the assignment problem would mean that the largest fraction is set to 1, and the other decision variables for a specific assignment are set to 0. This can result in a solution that does not satisfy all constraints anymore, which is less important in large, practical problems, where a very high accuracy is not required.

3.5.2 Mixed Integer Linear Programming

A Mixed Integer Linear Programming problem (MILP) is an LP where a subset of the decision variables is set to binary integer variables. Usually, an MILP is solved faster than an ILP, because an MILP is less constrained. An MILP can be solved with Lagrangian relaxation and branch-and-bound.

3.5.3 Dynamic Programming

Dynamic Programming (DP) is used by Kleywegt et al. (2004). Since the computation time of their solution is prohibitively high, it is not an alternative for our solution.

3.6 Main conclusions from the literature

The literature provides us with five main insights:

- (1) **A decomposed approach, which separates the planning procedure from the scheduling procedure, is better than an integrated approach.** The integrated approach tries to combine planning and scheduling in one step, meaning the timing and size of the deliveries is combined with the creation of trips for all deliveries. This increases the problem's complexity and uses a prohibitively large computation time.
- (2) **The long-term effect of a short-term decision can best be reflected by *relative delivery size*,** which is the size of the delivery to a customer divided by the capacity of this customer. Relative delivery size illustrates the 'earliness' of an order, the lower it is, the more the delivery is pulled forward. Since pulling an order forward is costly, because it increases the number of visits to a customer, the relative delivery size should be maximized.
- (3) **The short-term costs are mainly driven by the distance that needs to be travelled.** This distance can be expressed in three ways: Euclidean distance, real distance, and detour distance.
- (4) **To have an equal workload balance on the short-term, we should balance the total daily delivery volume in the planning period.** A different balancing criterion, such as working time or travelled distance, requires an integrated approach, and a calculation of routes together with the timing of the order. Since we use a decomposed approach, we balance on volume.
- (5) **To mitigate the seasonal peak in demand, we balance the workload over a longer period than just one week.** We take into account the future demand in this long-term workload balance, to stock customer such that they do not require a delivery during the peak in demand. In this way, we anticipate on knowledge about future peaks in demand.

4 Solution

Based on Chapters 2 and 3, we design an algorithm in this chapter, which is the solution to our problem. Paragraph 4.1 summarizes the main considerations in designing the algorithm and Paragraph 4.2 illustrates the complete algorithm. Paragraph 4.3 displays the alternative designs of the algorithm. We discuss the distance types and settings of the algorithm in Paragraphs 4.4 and 4.5.

4.1 Algorithm design

We summarize the main insights drawn from Chapters 2 and 3, which show that the design has four important objectives:

- (1) Create an efficient plan on the short-term, maximize the *volume per kilometre*.
- (2) Create an efficient plan on the long-term, minimize the *number of visits*.
- (3) Balance the workload on the short-term, to use the available resources efficiently.
- (4) Balance the workload on the long-term, to mitigate the seasonal peak.

The volume per kilometre measures the efficiency of the plan on the short-term. The relative delivery is a measure for the long-term efficiency, the higher the relative delivery, the fewer the number of visits to this customer. The workload can be balanced on the short-term by equalizing the volume that has to be delivered per day. An equal workload balance on the long-term is achieved by calculating the average demand per period for a number of periods into the future, and delivering that volume in the current planning period. Peaks in demand are foreseen, and customers are replenished before this peak in demand.

Paragraph 4.1.1 and Paragraph 4.1.2 give the solution requirements and the assumptions that are used as input for the design of the algorithm.

4.1.1 Solution requirements

- (1) Relatively low computation times for real-life instances.
- (2) The workload is balanced to cope with volatility on the short-term and the long-term.
- (3) Travel to different areas in the region on a single day to be able to respond to emergency orders.
- (4) Multi-depot approach, since a planner is responsible for a region with multiple depots in gas distribution.
- (5) Heterogeneous fleet approach, because there are different types and sizes of vehicles in real-life instances of the IRP in gas distribution.
- (6) The solution is implemented and tested in the ORION and SHORTREC suites.

4.1.2 Assumptions

- (1) Holding costs at the depot and at the customer are not taken into account.
- (2) Depots have an infinite supply of gas.
- (3) Depots are not assigned to a customer; the algorithm is free to choose by which depot a customer is served.
- (4) Customer usage is considered deterministic.
- (5) There is only one bulk product.

4.2 Algorithm description

We summarize the algorithm and illustrate the motivation behind the separate steps in Paragraph 4.2.1 and Figure 4.1. Paragraph 4.2.2 describes the algorithm in detail.

4.2.1 Algorithm overview

We start the algorithm with a set of customers, with certain start inventories and historical data, to calculate future usage. We finish the algorithm with a set of planned orders for a selection of these customers, on specific days in the planning period.

Step 1: Customer selection

We select the set of customers for the algorithm. We only select those customers that can at least receive a relative delivery size of α .

Motivation: To minimize the long-term cost, we have to minimize the number of visits to each customer. A high relative delivery size decreases the number of visits, since the optimal relative delivery size is 100% (Dror and Ball, 1987; Golden et al., 1984).

Step 2: Flexible clustering

We create clusters based on the set of ‘must-go’ customers.

Motivation: To minimize short-term cost, we have to minimize the distance travelled. By creating clusters based on distance, we aim to minimize the travelled distance (Campbell and Savelsbergh, 2004; Jung and Mathur, 2007).

Step 3: Generate schedules

We generate all possible delivery schedules for the customers that are in a cluster. A delivery schedule contains the relative delivery sizes per day of the planning period. For ‘may-go’ customers, it also contains a relative delivery size that corresponds with the end of the following planning period.

Motivation: The ILPs in Steps 4 and 5 need these schedules as input.

Step 4: Assignment of seed customers with an ILP

An ILP assigns seed customers to delivery schedules. The ILP plans all seed customer deliveries on a specific day and balances the number of seed customer deliveries over the depots and the planning period, with regards to the available vehicle capacity.

Motivation: The ILP has the objective to minimize the long-term cost by maximizing the relative delivery size (Dror and Ball, 1987; Bard et al., 1998). The number of seed deliveries is balanced, and not the volume, since we want to spread out the number of clusters, and thus the number of seeds, over the planning period. The volume is balanced in Step 5.

Step 5: Assignment of clustered customers with an ILP

An ILP assigns clustered customers to delivery schedules. The ILP plans all customer deliveries on a specific day in this step and balances the total volume of seed customer deliveries and clustered customer deliveries over the planning period, with regards to the available vehicle capacity.

Motivation: Balancing the workload is done by balancing the delivery volume per day in the planning period (Bard et al., 1998). We calculate the average demand for a number of planning periods into the future to balance the workload over a longer period. Thereby, the algorithm mitigates the seasonal peak in customer deliveries.

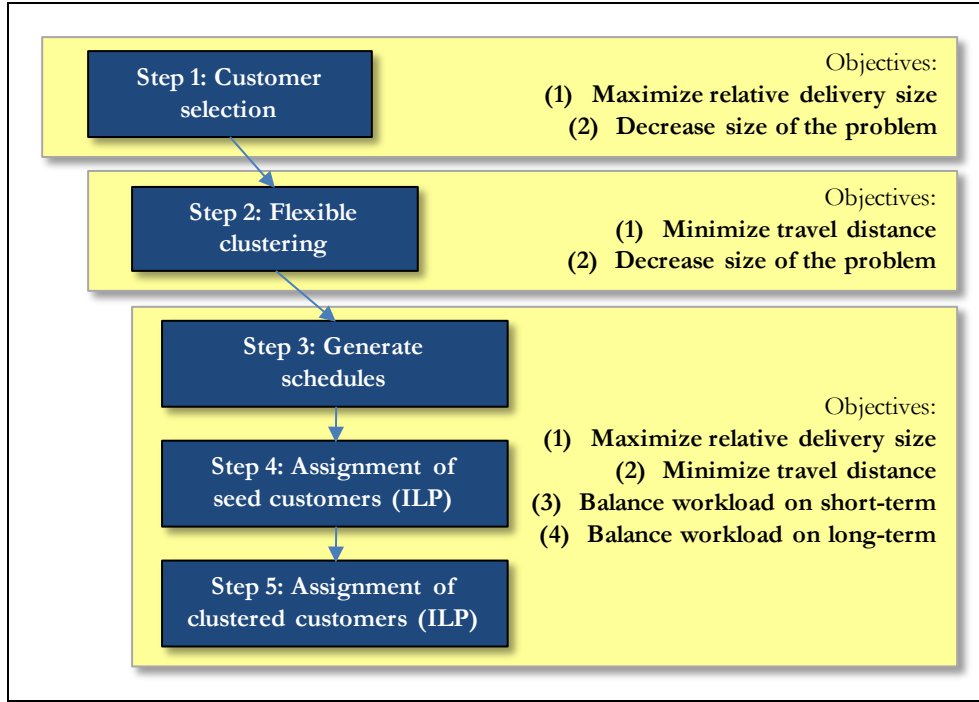


Figure 4.1. An overview of the steps in the algorithm and their objectives.

4.2.2 The algorithm in detail

This paragraph explains the details of the separate steps in the algorithm and presents mathematical models that are used. Appendix E illustrates all mathematical symbols that are used in this chapter.

4.2.2.1 Step 1: Customer selection

The driver for long-term costs in the IRP is relative delivery size (Golden et al., 1984; Dror and Ball, 1987). The relative delivery size is given in the left term of Equation 4.1, and to maximize it, we choose to only consider customers that can receive a relative delivery above α . This is the approach of Golden et al. (1984).

$$\frac{\text{Possible Delivery Volume}}{\text{Capacity}} \geq \alpha \quad (4.1)$$

In Equation 4.1, maximum stock is the capacity of the customer's tank that can be used because of safety reasons and regulations, and it is 85% of the total capacity of the customer's tank. We can not set α too high, since it would result in stock outs.

Must-go' customers

Additionally, all customers that need to receive a delivery in the planning period are added to the set of selected customers. This group holds forecasted 'must-go' customer that are pushed and in the VMI-program, and customers with pull- or call-in orders, that place the order themselves. Pull-orders are delivered within a time period after the call, which is agreed in a contract. We add these pull orders to the set of 'must-go' customers.

4.2.2.2 Step 2: Flexible clustering

In order to minimize travel distance, clustering is a tool that is widely applied in the literature (Campbell and Savelsbergh, 2004; Jung and Mathur, 2007). Clustering is done based on distance between a customer and a certain reference point, which is called a *seed*.

In the literature, clustering is usually done before the actual planning phase and is no daily activity of a planner. In a large customer set, information on the customer base changes continuously, thus to ensure that clusters are up to date, clustering should be part of the planning process. In our solution, clusters are created in a fast procedure, by selecting a seed and adding other customers to this seed's cluster. The closest customer is added first, and we add customers until the distance between a customer and the seed is above a certain threshold. All customers that are in the selection of customers can be selected into the clusters, including 'must-go' customers. These *flexible clusters* are computationally fast and based on the latest knowledge about the customer base.

The 'must-go' customers form the foundation of the problem (Campbell and Savelsbergh, 2004) and are therefore the basis for creating the clusters. The 'must-go' customer list is sorted, and the first customer on the sorted 'must-go' customer list is selected as a seed for the first cluster. We proceed until we have clustered all 'must-go' customers, either as a seed for a cluster or a customer that is selected in a cluster. A cluster can not include customers that were clustered earlier.

Sorting the 'must-go' customer list

The 'must-go' customer list is sorted on three criteria, which are explained in detail later:

- (1) The number of days between the start of the planning period and the latest day of delivery of the customer: in ascending order.
- (2) The number of deliveries for a customer in the planning period: in descending order.
- (3) The number of vehicle types that can visit the customer: in ascending order.

Rank	Customer	Days between start of planning period and delivery	Number of deliveries in planning period	Number of vehicle types
1	Customer A	3	2	1
2	Customer B	3	2	>1
3	Customer C	3	1	>1
4	Customer D	4	1	1
5	Customer E	4	1	>1

Table 4.1. An example of a sorted 'must-go' customer list.

Table 4.1 illustrates a sorted 'must-go' customer list with five 'must-go' customers. Customers on the 'must-go' customer list can also be selected in a cluster. To optimize the use of clusters, we do not want a customer in a cluster to have a latest delivery date *earlier* than the seed of that cluster, because that would cause problems in assigning the orders to days. The seed customers are assigned first, and if the seed customer is assigned to a later day than the latest delivery day of another customer in that cluster, we would lose the added value of clustering, since the other customer can not be assigned to the same day as the seed customer anymore.

A customer that requires multiple deliveries in the planning horizon should be a seed customer, since the size of its cluster is larger to be able to deal with the multiple deliveries. This is explained in the part on cluster constraints below.

Additionally, we want to select the most restricted customers as a seed first, because it makes it easier to combine several ‘must-go’ customers into one cluster. All customers can be added to a seed that can be visited by only one vehicle type, and fewer customers can be added to a seed that can be visited by all vehicles. This is explained in the part on restrictions below.

Restrictions

In Paragraph 2.2, we pointed out that vehicle and equipment restrictions are specific to gas distribution. For both of these restrictions there is a simple rule to follow in creating the flexible clusters. The seed of a cluster is the most important customer in the cluster, since it is the ‘must-go’ customer on which the cluster is founded. To make sure that a complete cluster can be visited by one vehicle, we need to make sure the vehicle that will visit the seed customer can also visit all the other customers in the cluster. This simple rule makes it easy to evaluate if a certain customer can be added to a cluster.

Cluster constraints

Constraint parameters for the clusters are used. A maximum distance between the seed and the customer will ensure proximity of clustered customers to the seed customer. A minimum and maximum number of customers in one cluster are used to control the size of the cluster. This minimum and maximum number of customers is multiplied by the number of required deliveries of the seed customer to ensure the cluster is large enough.

Number of seeds per day

Additionally, we ensure that every day has at least one or more seeds, by first generating one or more clusters for every day in the planning period. We select the first ‘must-go’ customer on the list which has a latest delivery day that matches the day in the planning period, for which we are currently selecting a seed customer. We proceed with the next day in the planning period with the same procedure. If there are no ‘must-go’ customers left for a certain day, a ‘may-go’ customer may be selected as a seed customer. It is beneficial to find multiple seeds per day in the planning period, since this geographically spreads the customers throughout the region, which is good to handle emergency deliveries, and which helps in balancing the workload throughout the region for the depots.

4.2.2.3 Step 3: Generation of schedules

After clustering, there are seed customers and clustered customers. We have selected these customers on relative delivery size and on distance to the seed customers. All customers that are not selected or not clustered are not considered further to keep the problem small and tractable.

To plan orders for the seed and clustered customers, we have to generate schedules which reflect the evolution of demand throughout the planning period. These schedules give the relative delivery size on a day in the planning period and an example is given in Table 4.2. The schedules F111 are for a ‘must-go’ customer requiring two deliveries, where the schedule with index 1 has a delivery of 75% of the customer’s maximum capacity on Monday and 35% on Tuesday. The schedules F222 are for a ‘must-go’ customer requiring one delivery, but with no delivery window on Wednesday.

The schedules F333 are for a ‘may-go’ customer with no delivery window on Sunday, and the 86% in the schedule with index 7 indicates that if the customer is not visited in this planning period, the customer requires a delivery of 86% of its maximum capacity by the end of the next planning period.

Schedule	Index	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday	End of next planning period
F111	1	75%	35%	0	0	0	0	0	0
F111	2	75%	0	50%	0	0	0	0	0
F111	3	75%	0	0	65%	0	0	0	0
F222	1	60%	0	0	0	0	0	0	0
F222	2	0	65%	0	0	0	0	0	0
F222	3	0	0	0	75%	0	0	0	0
F333	1	50%	0	0	0	0	0	0	0
F333	2	0	53%	0	0	0	0	0	0
F333	3	0	0	56%	0	0	0	0	0
F333	4	0	0	0	59%	0	0	0	0
F333	5	0	0	0	0	62%	0	0	0
F333	6	0	0	0	0	0	65%	0	0
F333	7	0	0	0	0	0	0	0	86%

Table 4.2. An example with 13 schedules illustrating the delivery percentages per day in the planning period.

4.2.2.4 Step 4: Schedule assignment with an ILP for seed customers

Seed customers are the distance reference points for assigning the clustered customers in Step 5, thus we assign the seed customers to schedules first.

An ILP is used to assign the seed customers to a schedule. The schedules are connected to a certain day of delivery, or when there are multiple deliveries for a customer in the planning period, to certain days of the planning period. A seed customer is nearly always a ‘must-go’ customer, and a seed customer can not be forwarded to a following period, since we need the seed customers as a distance reference for the clustered customers.

We want to balance the number of clusters per day and per depot throughout the planning period, to ensure efficient use of capacity. Therefore we balance the number of seed deliveries, and not their delivery volume. We use seed deliveries, and not seed customers, since a seed customer may have multiple deliveries in the planning period.

Workload calculation

Appendix F explains these procedures in detail, but we explain the concepts here. We balance the number of seed deliveries over the planning period and over the depots, and the maximum workload is calculated by multiplying the total number of seed deliveries with the fraction of the total available vehicle capacity in the planning period, that is available at the specific depot, on the specific day. We select the maximum of this workload, and the expected number of seed deliveries for a depot and a day to ensure the problem is feasible. We can add a certain percentage of the maximum workload that can be used as a bandwidth for feasibility, but we do not use it in this part of the thesis.

The use of a minimum workload for an ILP will result in an even balance of the number of seed deliveries over the week, since it forces the number of seed deliveries to be at least a certain number of deliveries on a specific day. We calculate a minimum workload in the same way as the maximum workload.

Multiple deliveries

Multiple deliveries are taken into account in the approach, since the total number of seed deliveries is used as input for the workload constraints and every day of delivery in a delivery schedule is considered.

Vehicle capacity

If there is no vehicle capacity available at any of the depots that can serve customer c on the day of delivery in schedule s , the combination of schedule s and customer c is excluded from the analysis. If there is no possible schedule left for customer c , the schedule will be allowed to ensure the ILP is feasible.

Figure 4.2 illustrates the ILP to assign seed customers to schedules. The objective is to maximize the total relative delivery volume, and to use the distance to the depots to balance the workload over the depots.

$$\begin{aligned} \max \quad & \sum_{p \in D} \sum_{c \in SC} \sum_{s \in S_c} \sum_{t=1}^T \left(y_{ts} + \frac{100}{d_{cp}} \right) X_{scp} \quad (4.2) \\ \text{s.t.} \quad & \sum_{p \in D} \sum_{s \in S_c} X_{scp} = 1 \quad \text{where } c \in SC \quad (4.3) \\ & w_{\min tp} \leq \sum_{c \in SC} \sum_{s \in S_c} m_{st} X_{scp} \leq w_{\max tp} \quad \text{where } t = 1, \dots, T, p \in D \quad (4.4) \\ & X_{scp} \in \{0, 1\} \quad \text{where } s \in S_c, c \in SC, p \in D \quad (4.5) \end{aligned}$$

Variables

X_{scp} binary variable that equals 1 if delivery schedule s is assigned to customer c , and depot p

Data

y_{ts} relative delivery size on day t in schedule s (%)

d_{cp} distance contribution to assign customer c to depot p

$w_{\max tp}$ maximum number of seed deliveries on day t , depot p

$w_{\min tp}$ minimum number of seed deliveries on day t , depot p

m_{st} 1 if there is a delivery on day t in schedule s , 0 otherwise

S_c set of allowed delivery schedules for customer c

SC set of seed customers

D set of depots

T number of days in the planning horizon

Figure 4.2. The ILP to assign the seed customers to delivery schedules balances the number of seeds over the days in the planning period.

The ILP has the objective to maximize the relative delivery size to the seed customers (Dror and Ball, 1987). Additionally, the distance to the depots is included in the objective function to balance the number of seed deliveries over the depots and to choose the closest depot for each seed. These two measures are incorporated into one objective function, which is tested in practice and works best in this form.

Constraint 4.3 is added to ensure that one of the delivery schedules is chosen, because we need the seed deliveries as distance reference points for Step 5. Constraint 4.4 requires that the number of seeds assigned to a day and a depot is between a certain minimum and maximum workload. Constraint 4.5 outlines the binary property of the decision variables. Since it is a small problem instance, branch-and-bound is fast and can be used.

4.2.2.5 Step 5: Schedule assignment with an ILP for clustered customers

We have assigned the seed customers to delivery schedules in Step 4. Thus in this step, we use the knowledge on the planning of the seed deliveries as a distance reference.

We assign the clustered customers with an ILP, and the objective of this step is to minimize the distance travelled, by minimizing the distance between a clustered customer and a seed delivery. Therefore, we assign a clustered customer to a delivery schedule *and to a seed delivery*, which increases the number of variables.

We create a proper workload balance for the planning period, by balancing on delivery volume. The workload is restricted by a minimum and maximum workload per seed delivery, and the balance over the planning period depends on the available vehicle capacity at a depot on a certain day. The workload is calculated to balance the delivery volume not only over the week, but also over a longer period. This is done by making use of the forecasted demand in the longer period. We explain this later in this paragraph.

Costs for excluding ‘may-go’ customers from the planning period

‘May-go’ customers can be forwarded to a future planning period, and they have a delivery schedule that represents this forwarding decision. The costs for excluding customers from the planning period are used to determine which ‘may-go’ customers should be planned in the current planning period. Equation 4.6 illustrates the forwarding cost function for ‘may-go’ customers. This expression is also used in the Period Scheduler.

$$\frac{\tilde{y}_c^3}{10^4 \cdot \sqrt[4]{\tilde{d}_c}} \quad \text{where } c \in CC \quad (4.6)$$

\tilde{y}_c is the relative delivery size when customer c is not delivered in this planning period, this is the relative delivery size at the end of the following planning period. \tilde{d}_c is the distance from customer c to its nearest seed delivery. CC is the set of clustered customers. Equation 4.6 is based on the notion that a customer with a higher relative delivery next week is more urgent, and thus more interesting to include in the current planning period. Additionally, a customer that is relatively close to a seed customer in the current planning period is more interesting to include as well, since the increase in total travel distance when the customer is added is low. The term offers an exponential growth in costs if the relative delivery decreases, or the distance to the nearest seed delivery increases.

Workload calculation

The total delivery volume for the week is determined by calculating the average demand per planning period for the next n planning periods. This is used in the Period Scheduler as well, to cope with peaks in future demand and to ensure that customers have sufficient stock so they require fewer deliveries when the real peak in demand is there (Welch et al.,

1971). This total delivery volume needs to be adjusted if the volume for ‘must-go’ customers is higher than the calculated average demand per planning period. This may happen when the demand is decreasing in the next n weeks.

The workload is balanced evenly over the depots and the days in the planning period, depending on the available vehicle capacity at a depot on the specific day in the planning period. We have already balanced the seed customer deliveries with the same philosophy. We calculate the workload per depot and per day in the planning period by multiplying the total delivery volume for the week with the fraction of the total available vehicle capacity in the planning period that is available at the specific depot, on the specific day. We select the maximum of this workload, and the expected volume per depot and day, to ensure the problem is feasible. We add a certain percentage of the maximum workload that is used as a bandwidth β for feasibility.

Now that we have calculated the workload per depot and per day, we can easily transfer this to a workload per seed delivery, since we know the assignment of the seed delivery to depots and days. We divide the workload per depot and day by the number of seed deliveries assigned to this depot on this day. Additionally, we subtract the delivery volume of the specific seed delivery to get the workload per seed delivery.

Next to a maximum workload, a minimum workload is added to the ILP to find a well-balanced plan. The calculation of the minimum workload is based on the same principles as the maximum workload, but the bandwidth β is now subtracted. Appendix G explains these procedures in detail.

Multiple deliveries

Multiple deliveries in one week are rare in the gas industry, but do occur. Customers with multiple deliveries will not be in the clustered set, but in the seed customer set. This is because the clustering is done in a specific order, in which a customer with multiple deliveries is chosen as a seed customer before a customer with a single delivery. If it does occur that a customer with multiple deliveries is in the clustered set, we add all deliveries after the first delivery to the closest seed on the specific day of the delivery. Thus the distance contribution of assigning a customer to a seed delivery and schedule is raised with the additional distance for the deliveries after the first delivery in the schedule. The volume that is expected to be delivered on that specific day in the assigned schedule is also added as a volume contribution for that nearest seed delivery.

Vehicle capacity

If there is no vehicle available at any of the depots that can serve customer c on the day of delivery in schedule s , the combination of schedule s and customer c is excluded from the analysis. If there is no other possible schedule left for customer c , the schedule will be allowed to ensure the ILP is feasible. The customer is then assigned to a day where there is no vehicle capacity for the customer.

Figure 4.3 illustrates the ILP to assign clustered customers to schedules.

$$\min \sum_{k^p \in SCD} \sum_{c \in CC} \left(\left(\sum_{s \in S_c} d_{ck^p} \cdot X_{sck^p} \right) + \left(\frac{\tilde{y}_c^3}{10^4 \cdot \sqrt[4]{\tilde{d}_c}} \right) X_c \right) \quad (4.7)$$

s.t.

$$\sum_{k^p \in SCD} \sum_{s \in S_c} X_{sck^p} + X_c = 1 \quad \text{where } c \in CC \quad (4.8)$$

$$v_{\min k^p} \leq \sum_{c \in CC} \sum_{s \in S_c} r_{st} X_{sck^p} + \sum_{c \in CC} \sum_{s \in S_c} \sum_{g \in SCD} R_{sk^p} X_{scg} \leq v_{\max k^p} \quad (4.9)$$

where $t = 1, \dots, T$; $k^p \in SCD$

$$X_{sck^p} \leq r_{st} \quad \text{where } t = 1, \dots, T, k^p \in SCD, c \in CC \quad (4.10)$$

$$X_c = 0 \quad \text{where } c \in CC_{mustgo} \quad (4.11)$$

$$X_{sck^p} \in \{0,1\} \quad \text{where } s \in S_c, c \in CC, k^p \in SCD \quad (4.12)$$

$$X_c \in \{0,1\} \quad \text{where } c \in CC \quad (4.13)$$

Variables

X_{sck^p} binary variable that equals 1 if delivery schedule s is assigned to customer c , and seed delivery k^p , 0 otherwise

X_c binary variable that equals 1 if customer c is excluded from the current planning horizon, 0 otherwise

Data

\tilde{y}_c relative delivery size if customer c is forwarded, the relative delivery size at the end of the next planning period (%)

d_{ck^p} distance contribution to assign customer c to seed delivery k^p

\tilde{d}_c is the distance from customer c to its nearest seed delivery

k^p seed delivery k that is assigned to day t , and depot p

$v_{\max k^p}$ maximum workload for seed delivery k^p

$v_{\min k^p}$ minimum workload for seed delivery k^p

r_{st} integer variable for the delivery volume on day t in schedule s

S_c set of allowable delivery schedules for customer c

R_{sk^p} volume for the deliveries after the first delivery in schedule s that is assigned to seed k^p

SCD set of seed customer deliveries

CC set of clustered customers

CC_{mustgo} set of clustered customers that are ‘must-go’ customers

T number of days in the planning horizon

Figure 4.3. The ILP to assign the clustered customers to schedules balances the delivery volume over the planning period.

The ILP has the objective to minimize the distance between the seed and the customer. An expression is added for the ‘may-go’ customers, since they have the option to be excluded from the planning period, which comes with certain costs. Constraint 4.8 is added to ensure that only one schedule is chosen. Constraint 4.9 ensures that the total volume that is assigned to a seed customer delivery is between the minimum and maximum workload. Because we do not want a schedule s to be assigned to a seed customer delivery k^p if there is no delivery to the customer on day t in that schedule, Constraint 4.10 is added. Constraint 4.11 ensures that a ‘must-go’ customer can not be forwarded to a future planning period. Constraints 4.12 and 4.13 outline the binary property of the two decision variables. We round the LP relaxation of this ILP to solve it, since the number of decision variables is relatively high.

Scheduling

After assigning every customer to a schedule, the day for every customer delivery is calculated. The next step is to schedule these orders with SHORTREC. SHORTREC finds a trip schedule for every day in the planning period. One can decide to eventually only use the schedule for the first day. In this way, we get a rolling horizon, similar to the solution of Bard et al. (1998), in which the algorithm is performed every day, but the customer deliveries are balanced over the long-term and the short-term.

4.3 Alternative algorithm designs

Five alternative algorithm designs are given, and they are tested in Chapter 5.

Balance customers

Campbell and Savelsbergh (2004) state that a vehicle should always leave the depot filled up completely, but due to route failures, the exact delivery size of a trip is not known in advance (Dror and Trudeau, 1988). By adding a balance customer at the end of every trip with a flexible delivery volume, we increase utilization and we have a higher probability of emptying the vehicle before returning to the depot. Since a balance customer is a ‘may-go’ customer, there is no real problem if the delivery can not be made. The application of this concept is explained in Appendix I.

Alternative objective functions

We have created six different designs for the objective function in the ILP for assigning clustered customers to schedules and seed customer deliveries. The six different objective functions are given in Appendix J.

Check for customer combinations

Since distance should be minimized, we ensure that the vehicle that visits the seed customer may also visit all other customers in a cluster. Therefore, we check if a customer can be delivered by the same type of vehicle as the seed customer before assigning the customer to the seed customer. Appendix K illustrates the ILP for clustered customers in this case.

Vehicle type knowledge

Additionally, we can make use of the knowledge about vehicle type restrictions per customer. Thus, if a customer can be visited by a single type of vehicle, we are certain that this type of vehicle is needed for this customer. If a customer can be visited by more than one vehicle, this certainty is lost, since there are multiple options, and therefore, we

can only separate between groups of customers that can be visited by one vehicle type and customers that can be visited by more than one vehicle type.

We use the knowledge about the capacity of these vehicles to determine the total amount of workload that can be assigned to the seed customer. By doing so, the workload for a vehicle type is more equally divided over the week and over the depots.

Period Scheduler seed selection procedure

We test the seed selection procedure of the Period Scheduler, explained in Paragraph 2.4 and Appendix A. The seed selection procedure selects as many seed customers around a depot as there are days in the planning period. The seed selection procedure replaces Step 2, the flexible clustering in the algorithm. The seed customers are assigned to days and depots in Step 4, and not assigned to days randomly like in the original Period Scheduler seeds selection procedure. Therefore, we have no certainty that there is an equal balance of clusters in the planning period.

4.4 Distance types

The several distance types that can be used in the tests are the following distance types:

- (1) Euclidean distance, distance that is given by measuring a straight line on a map between two locations.
- (2) Real distance, distance that is given by measuring the minimum travel distance through a given network between two locations.
 - a. The distance measured in travelled kilometres.
 - b. The distance measured in travelled time, by computing the expected travel time between two locations.
- (3) Detour distance, distance that is given by the additional real travel distance when a certain customer is inserted in an existing trip between a seed order and the depot.

Detour distance can not be used for the ILP for seed customers, since there is no other customer to use as a reference. Table 4.3 shows the distance types applicable to the ILPs.

ILP	Euclidean	Real kilometres	Real time	Detour
Step 4: ILP for seed customers	✓	✓	✓	
Step 5: ILP for clustered customers	✓	✓	✓	✓

Table 4.3. The distance types that can be used for the ILPs.

4.5 Parameters of the algorithm

Table 4.4 displays the steps and parameters of the algorithm.

Algorithm	Parameters
Step 1: Customer selection	α, T
Step 2: Flexible clustering	minimum distance, minimum cluster size, maximum cluster size, minimum number of seeds per day, distance type
Step 3: Generate schedules	-
Step 4: ILP for seed customers	distance type
Step 5: ILP for clustered customers	β (bandwidth clustered customers), n (workload is average of n planning horizons of forecasted demand), distance type

Table 4.4. The algorithm, and the parameters that can be set at each specific step.

5 Computational results

The algorithm is tested through a series of steps:

- (1) Paragraph 5.1 explains the computational experiments and analyses the dataset.
- (2) Paragraph 5.2 compares the results of the algorithm with what is planned in practice.
- (3) Paragraph 5.3 tests the sensitivity of several settings in the algorithm.
- (4) Paragraph 5.4 tests the alternative designs of the algorithm.

We only compare the algorithm's results and the base scenario's results on a weekly basis; the algorithm plan of a week is not fed back into ORION to use as a basis for the next week. We could not develop a test version that copes with feedback loops in this research. The most important measure on the short-term is the *volume per kilometre*, and the most important measure on the long-term is the *average relative delivery size*, which is the average relative delivery size for all planned orders.

5.1 Experimental design

For the experiments, we make use of actual data. The customers, usage curves, and inventory levels are obtained from a database that also stores actual deliveries. We compare the planning results of the algorithm with the actual deliveries to see if there is any improvement. We can not compare to the originally planned deliveries, since they were overwritten by the actual deliveries in the database. An Intel Pentium 4 computer with a 3.0 GHz CPU and 2.0 GB internal memory (RAM) is used for the experiments. This paragraph discusses the algorithm settings, the dataset, and the general assumptions.

5.1.1 Algorithm settings

Appendix L illustrates the settings used to compare the results for configuration of the model. Some of these settings, such as the number of seeds per day, were found by tests that are not illustrated here.

5.1.2 Dataset

The data are taken from a real-life dataset and concerns a region with three depots in the Northern-England region (NE-region) depicted in Figure 5.1. These three depots serve around 3.100 customers, and distribute around 12,8 million liters annually. The number of vehicles is different throughout the year, but generally there are two vehicles per depot per day. On Saturdays there are fewer vehicles available, and on Sundays there are usually no vehicles available. This fluctuation is a result of the volatility of gas demand, and it requires that vehicles are available on Sunday in the winter for instance. The vehicle capacity is also different per vehicle. For an overview of vehicles with their capacity, see Appendix M.

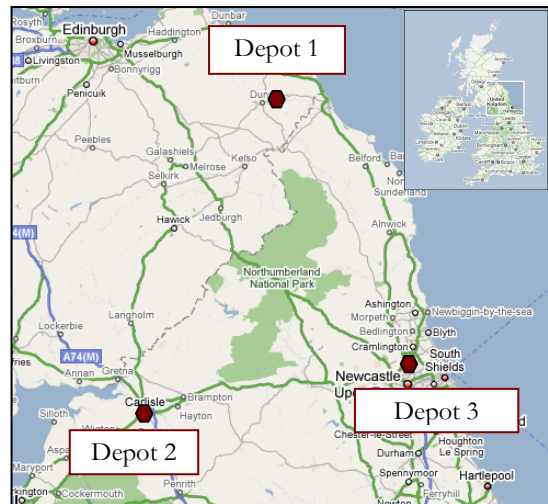


Figure 5.1. The location of the depots, and an illustration of the NE-region.

60% of the customers are in the VMI-program, and their demand is forecasted by ORION. For the customers that are not in the VMI-program demand is not forecasted, and we expect that their demand is proportional to the demand of the VMI customers. Table 5.1 illustrates the data for VMI-, and non-VMI-customers in the NE-region.

Customer type	Total demand in 2007 (litres)	Relative
VMI customers	12.069.935	94,35%
Non-VMI customers	723.342	5,65%
Total	12.793.277	100,00%

Table 5.1. Exactly 94,35% of the total demand is from VMI-customers. This is used in calculating the expected total average demand for a specific planning period.

To find the total expected demand, we assume that the seasonal pattern of the non-VMI customers is equal to the VMI customers, and we divide the total expected forecasted demand by the percentage of volume for VMI-customers that is observed in historical data. Thus in the case of the NE-region, we divide a total expected forecasted demand of 12,1 million litres by 94,35%, and we get a total expected demand of 12,8 million litres.

On average, nine stops per trip occur, and a vehicle can perform multiple trips. Most customers can have a delivery every day of the week, except for Sundays. The delivery windows are from 6:00hrs in the morning until 18:00 hrs in the evening. The working time of the depots is between 6:00 hrs and 22:00 hrs. Drivers work between 6:00 hrs and 18:00 hrs, and they have a lunch break between 12:00 hrs and 13:00 hrs of 45 minutes.

The planners use ORION to generate orders, but they do not use SHORTREC to schedule routes. Their program does not consider any trip restrictions, so their routes can not be used for an equal comparison. Therefore, we copy all delivered orders, and only use the date of delivery. We schedule these orders with SHORTREC to find feasible and efficient routes. We only consider the customers that were originally assigned to the three depots and not the customers outside this region. We can compare our results to these routes, which we call the 'base scenario'. Appendix M illustrates a comparison between the schedules made by the planners and the schedules created with SHORTREC.

The data from weeks 40, 41, and 42 in the year 2007 are chosen, since these are busy weeks with many pull orders, just before the winter peak period. The data for other weeks (weeks 14-17, 27-30, and 48-50) can not be used, because the forecast data were adjusted in 2008. This means the forecasts are higher than originally used, and as a result, many inventory levels are already too low before the start of the week. This causes the algorithm to plan around 50% of the total volume for a planning period on the first day. This is very unrealistic, and therefore not a good representation of an actual week.

With the settings described in Paragraph 5.1.1, the total average demand per week, the volume of must-go orders, and the total volume of selected orders is given by Table 5.2.

Week	Expected average total weekly demand (litres)	Total 'must-go' volume (litres)	Total volume in customer selection (litres)	Total number of customers in clusters
40	307.660	190.885	434.456	480
41	312.649	184.174	433.727	466
42	321.215	161.519	450.962	505

Table 5.2. The volumes to calculate the workload for the ILP for clustered customers.

5.1.3 Assumptions

- (1) The relationship between the demand of non-VMI customers and VMI-customers is the same on every day of the year, and can be found by analyzing historical data.
- (2) Vehicle capacity is equal on every day of the week, except for Sundays when it is zero. We also test a set-up where Saturdays have 50% of the vehicle capacity in Paragraph 5.2.3.
- (3) Supply at depots is unlimited.
- (4) Driver overtime is not allowed.

5.2 Comparison base scenario and algorithm

This paragraph compares the performance measures of the algorithm plans with the actual plans. Paragraph 5.2.1 presents the comparison and illustrates the geographical results to show that the algorithm balances the workload over the depots. Paragraph 5.2.2 illustrates the workload balance on the short-term and the long-term. Paragraph 5.2.3 discusses the results when the vehicle capacity is differently set-up, to illustrate the responsiveness of the algorithm to available vehicle capacity.

5.2.1 Weekly comparison

Table 5.3 illustrates that the algorithm causes an increase in the average relative delivery, which results in fewer visits on the long-term. Additionally, there is an increase in volume per kilometre, which results in a more efficient plan on the short-term. The decrease in cost per litre concerns a decrease in fictive costs used by SHORTREC, and indicates an improvement in overall costs.

Scenario	Volume (litres)	Dist. (km)	Volume/ kilometre (litres/km)	Cost/ litre	Veh. used	Average relative delivery	Volume/ kilometre improve.	Cost/ litre savings
Base scenario	718.457	13.448	53,4248	0,0581	56	53,55%	-	-
Algorithm	981.942	15.095	65,0508	0,0507	68	59,77%	21,76%	12,85%

Table 5.3. The results illustrate that the algorithm performs better than the base scenario in weeks 40, 41, and 42. The separate week results are illustrated in Appendix O.

Geographical results

Table 5.4 illustrates that the workload is equally divided over the depots in the planning period, according to the available vehicle capacity.

Depot	Available vehicle capacity (%)	Planned volume algorithm (%)	Planned volume base scenario (%)
Depot 1	40,91	37,50	37,28
Depot 2	30,68	29,44	32,50
Depot 3	28,41	33,06	30,22

Table 5.4. The equal workload balance over the depots in the planning period.

In addition to the results in figures, the geographical results are important to determine the spread of the workload through the region. Figure 5.2 illustrates the geographical results, and one customer block can hold more customers of all different types (seed and clustered customers) that are in the same postcode area.



Figure 5.2. In the geographical spread of the algorithm results in week 42 can be observed that the workload is spread throughout the region.

The orders are evenly spread over the depots in the region, and on each day there is workload for all depots. The next paragraph illustrates the equal spread of workload over the planning period.

5.2.2 Workload balance

One of the objectives is to balance the workload on the long-term and short-term.

Long-term workload balance

The long-term workload balance is important in mitigating the seasonal winter peak. By setting the algorithm workload equal to the average demand for the next 13 weeks, a peak can be anticipated, and certain volumes can be delivered before the peak reaches its highest point. Since the algorithm plans a higher delivery volume in the early stages of the winter peak (weeks 40, 41, and 42), more volume is already delivered before the real busy times start. Additionally, the higher average relative delivery in the algorithm ensures that only the most urgent customers are visited. As a result, we ensure a lower stress on capacity in future periods, and lower volumes need to be delivered during the highest point of the winter peak.




Figure 5.3. The algorithm follows the 13 week average expected demand.

Figure 5.3 illustrates the difference in delivery volume if we use a 13 week average demand. Paragraph 5.3.2 describes the sensitivity of the number of weeks used for calculating the average.

Short-term workload balance

Figure 5.4 illustrates the workload balance in the planning periods for the algorithm plans. Figure 5.5 illustrates the plans of the base scenario. The Monday stands out in Figure 5.4, because the usage data were adjusted in the beginning of 2008, after the actual planning periods. The adjusted usage data influence the due dates of the customers, and for most customers, the due date is now earlier than the planning period. Therefore, all these customers are planned on the first day in the planning period. The workload balance of the algorithm for the other days is a lot better than the plans of the base scenario. This workload balance is influenced by the bandwidth β . The sensitivity of this parameter is tested in Paragraph 5.3.3.




Figure 5.4. The workload volume is evenly distributed by the algorithm.




Figure 5.5. The graph illustrates the workload balance of the base scenario.

The available capacity on Saturday is lower than during the week in the base scenario. Therefore, there is no planned volume on Saturday in week 41 and little volume in week 42. The available capacity during the other days of the week is equal. In Figure 5.5, Monday does not stand out like in Figure 5.4, because the planners used different usage curves since they were adjusted in 2008. Appendix P illustrates the working time balance for both the base scenario and the algorithm plans. The working time balance of the algorithm is also quite equal, which indicates there is a connection between the balance in delivery volume and the balance in working time.

5.2.3 Different vehicle capacity set-up

Figure 5.6 illustrates the responsiveness of the algorithm to the available vehicle capacity. Table 5.5 illustrates that the results are slightly better when the Saturday vehicle capacity for each depot is 50% of the vehicle capacity on any other day, which is a common setting in the NE-region. Appendix Q illustrates the weekly and the geographical results, which show that there is workload at every depot on every day of the planning period.

Scenario	Volume (litres)	Dist. (km)	Volume/ distance (litres/km)	Cost/ litre	Veh. used	Average relative delivery	Volume/ distance improve.	Cost/ volume savings
Base scenario	718.457	13.448	53,4248	0,0581	56	53,55%	-	-
Algorithm: Saturday 100%	981.942	15.095	65,0508	0,0507	68	59,77%	21,76%	12,85%
Algorithm: Saturday 50%	961.790	14.751	65,2017	0,0506	66	59,49%	22,04%	12,92%

Table 5.5. The results of running the algorithm with a different set-up of capacity.




Figure 5.6. The graph illustrates that the workload balance is adjusted according to the vehicle capacity adjustment.

5.3 Sensitivity analysis

We study the sensitivity of the main settings of the algorithm in this paragraph. Paragraph 5.3.1 describes the sensitivity of setting α in customer selection. Paragraph 5.3.2 illustrates the effects of different settings for the number of weeks we look into the future to calculate the average delivery volume. Paragraph 5.4.3 describes the sensitivity of setting the bandwidth that used to calculate the minimum and maximum workload for the ILP for clustered customers. We discuss the distance types in Paragraph 5.4.4.

Other settings, such as the number of seeds per day and the threshold distance in the clustering phase, are found by trial and error, and depend heavily on the geographical area for which the algorithm is used. The settings we discuss below have the highest impact on algorithm results.

5.3.1 Setting α in customer selection

We use the parameter α to select customers as the first step of the algorithm. The higher α , the lower the number of customers that is selected. Figure 5.7 shows the volume per kilometre and the average relative delivery as a function of α . Week 40 is not included, since the weeks 41 and 42 already illustrate the effects. Figure 5.8 displays the algorithm run time and the number of customers as a function of α . The algorithm run time does not include order and demand calculation in ORION, and scheduling of orders in SHORTREC. α is between 0 and 1, but most customers have a safety stock level of at least 20%, and a maximum capacity of 85%. Therefore, α leaves out all customers that are not emergency customers when α is set higher or equal to 0,765 $((85 - 20) / 85)$. Appendix R illustrates the separate results, as well as the number of customers in the selection, and the algorithm run times per setting of α .

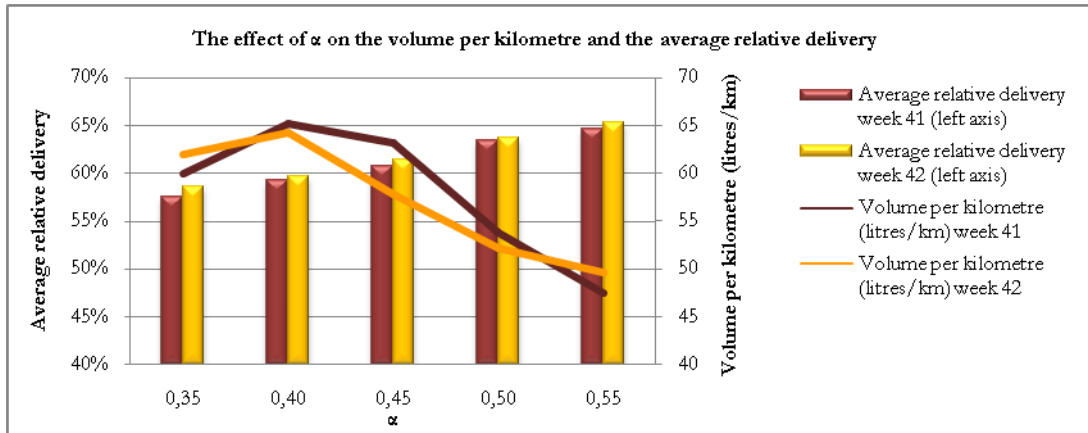


Figure 5.7. The graph illustrates the volume per kilometre (on the right axis) and the average relative delivery (on the left axis) as a function of α in the weeks 41 and 42.

The results for volume per kilometre are best if α is set to 0,40, but the long-term indicator of average relative delivery performs better with a higher setting for α . Additionally, a setting for α of 0,45 would diminish algorithm run time to below 600 seconds, or 10 minutes as can be seen in Figure 5.8. Therefore, an α between 0,40 and 0,45 is suitable for use when the algorithm is run frequently.

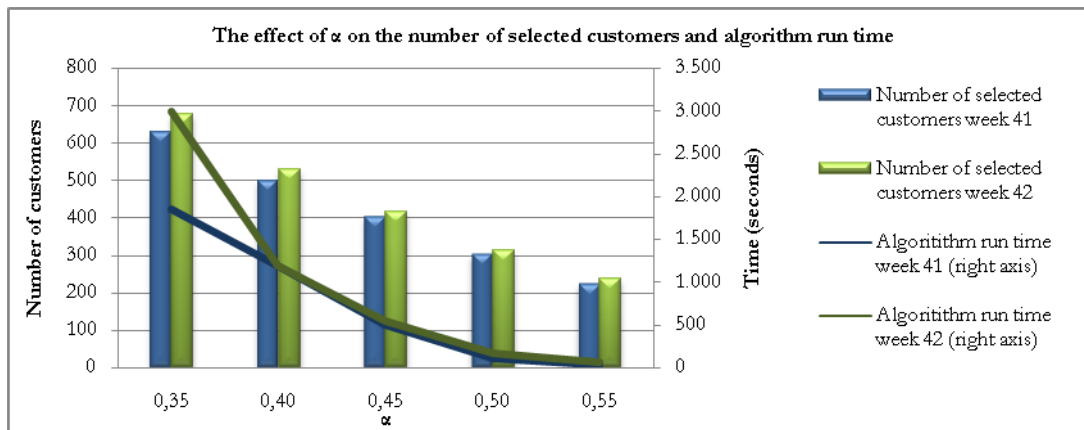


Figure 5.8. The graph illustrates the number of selected customers (on the left axis), and the algorithm run time (on the right axis) as a function of α in the weeks 41 and 42.

Since the demand pattern is not stable over time, α needs to be determined for multiple periods in the year. The problem in our dataset is the adjusted usage curves, which make it difficult to find an α that is generally applicable. Figure 5.9 illustrates this. We expect that the summer week selection (week 2007-27) holds fewer customers than the winter week selection (week 2007-50) with the same α , due to the higher demand in the winter period, but the opposite is true. The test is therefore inconclusive on which α to use in which part of the year.

For the period under observation (weeks 41 and 42) the number of customers for the week should be between 450 and 500, as we can see in Figure 5.8. An alternative to the use of α would be to select the best 400 ‘may-go’ customer with regards to relative average delivery, in addition to the ‘must-go’ customers. In that way, we do not need α as a separate setting anymore, and we have a fast algorithm which obtains good results.

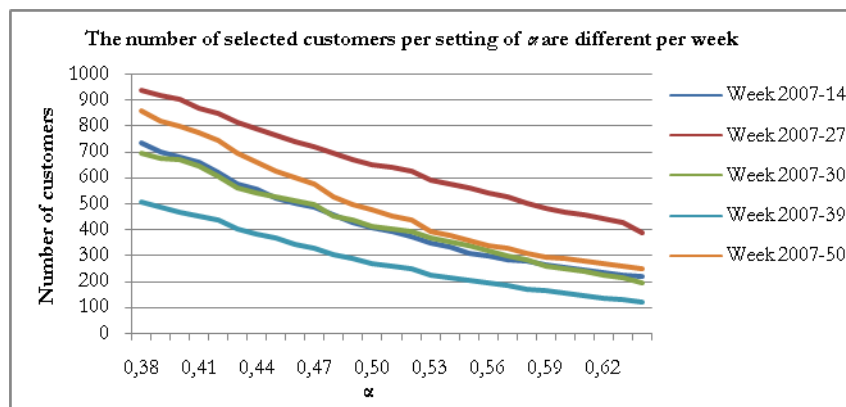


Figure 5.9. The graph illustrates the number of customers in the customer selection for different settings of α in different weeks of the year.

5.3.2 Long-term workload balance

The maximum workload for every planning period is calculated with an average demand n weeks into the future. The long-term workload balance is based on this concept, and the higher n , the flatter the curve of the delivery volume. Figure 5.10 illustrates the curves of the delivery volume with different values of n . The curve for 26 weeks stops at week 40, because there was insufficient data to calculate 26 weeks into the future beyond that point.




Figure 5.10. The graph illustrates the demand curves for n weeks into the future.

Figure 5.10 illustrates that, with the use of a 13 week average, the demand curve is still a steep curve, and that the peak reaches a high point even in the case of a 26 week average. Table 5.6 displays the results of using a 13 week average or a 26 week average. The 26 week average for week 49 (343.091 litres) is estimated by using the shape of the curves of the weekly average demand for lower n .

Week	Scenario	Average weekly demand (litres)	Volume (litres)	Average relative delivery
2007-17	Base scenario	-	258.581	54,23%
2007-17	13 week average*	190.507	267.900	71,04%
2007-17	26 week average*	181.323	267.900	71,04%
2007-30	Base scenario	-	236.659	49,52%
2007-30	13 week average*	172.137	259.541	65,77%
2007-30	26 week average	251.057	280.466	65,08%
2007-40	Base scenario	-	257.604	53,29%
2007-40	13 week average	307.660	319.827	60,30%
2007-40	26 week average	324.272	327.952	60,11%
2007-49	Base scenario	-	388.162	57,51%
2007-49	13 week average	353.376	366.939	65,73%
2007-49	26 week average	343.091	356.694	65,92%

* 'Must-go' volume is larger than calculated average weekly demand, thus the delivered volume is much higher.

Table 5.6. The table illustrates the results of testing a 13 and 26 week average.

The results of using a 26 week average results in a higher volume to be planned for week 30 and 40, and therefore these weeks show a slight decrease in average relative delivery. Week 50 shows a decrease in planned volume when 26 weeks are used, and therefore an increase in average relative delivery. We can conclude that the application of the 26 week average works well for this dataset, since the average relative delivery does not decrease dramatically.

5.3.3 Short-term workload balance

By setting the bandwidth for the ILP for clustered customers (β), the short-term workload balance is influenced. A higher bandwidth increases the workload freedom, and thus the standard deviation of the delivery volume in the planning period. Figure 5.11 illustrates the results in standard deviation over the week as a function of the bandwidth setting. One curve holds the standard deviation including Monday, and the other excluding Monday. The delivery volume on Monday is significantly higher than on the other days, due to the adjusted usage curves. Appendix S displays the results for volume per kilometre and average relative delivery, which show almost no change if the bandwidth changes.




Figure 5.11. The graph shows the standard deviation in volume per day of the planning period for week 42.

By using the relaxation and a rounding procedure in stead of an exact ILP solution procedure, some decision variables can cause disturbances in the workload volume per day. The total number of customers that are rounded during the rounding procedure is 30 customers, or 6,38% of the total number of customers of 470. These 30 customers represent 14,08% of the total volume that is assigned in the ILP, which shows that especially the larger deliveries have to be rounded after solving the LP relaxation. Although the effect is not large, an improvement of the ILP solution method leads to an improvement in the control of the short-term workload balance.

5.3.4 Distance types

Several distance types for the ILP for clustered customers are tested. The results in Table 5.7 indicate that detour distance is better than the other distance types. Appendix T illustrates the separate results per week, and the geographical results for real distance in kilometres and detour distance. These results illustrate that the use of detour distance assigns more customers between a depot and a seed.

Distance type	Volume (litres)	Dist. (km)	Volume/ kilometre (litres/km)	Cost/ litre	Veh. used	Average relative delivery	Volume/ kilometre improve.	Cost/ volume savings
Base scenario*	718.457	13.448	53,4248	0,0581	56	53,55%	-	-
Real distance (km)	951.115	15.528	61,2516	0,0525	68	60,77%	14,65%	9,64%
Real distance (time)	1.018.789	19.540	52,1386	0,0609	78	59,26%	-2,41%	-4,70%
Euclidean distance	975.112	16.004	60,9293	0,0531	72	60,16%	14,05%	8,65%
Detour distance	981.942	15.095	65,0508	0,0507	68	59,77%	21,76%	12,85%

* The base scenario was planned by the planners, and therefore has no such setting.

Table 5.7. The results for the weeks 40, 41, and 42 illustrate that detour distance is most efficient on the short-term measure volume per kilometre.

5.4 Testing alternative algorithm designs

We discuss the results of testing the alternative algorithm designs. The separate tests are given in Appendix U.

Balance customers

The tests illustrate that the use of balance customer is inefficient, since the average relative delivery and the volume per kilometre are lower. The average relative delivery is higher if no balance customers are used, since balance customers can receive a delivery that is lower than their maximum delivery. The vehicle utilization is higher when balance customers are used, because we maximize the balance customers, and fill up the vehicles after scheduling.

The use of balance customers is a more practical planning solution to cope with route failures, and its effect is better measured in a stochastic environment, where actual delivery volumes can be simulated. The extension can also be used for marketing or implementation purposes, since it is comparable to planning methods used in practice.

Objective functions

We test the different objective functions mentioned in Appendix J. The objective function used in the ILP in Step 5 of the algorithm performs best according to the tests in Appendix U.

Check for customer combinations

The results in Appendix U for using a customer combination check illustrate that it is better to leave the additional constraint out. Therefore, this functionality is not used in the algorithm.

Vehicle type knowledge

Paragraph 4.3 explains the use of knowledge about vehicle type restrictions to determine the total workload. The results in Appendix U illustrate that we obtain worse results when we use this functionality.

Period Scheduler seed selection method

We test the seed selection method used in the Period Scheduler. The seed selection procedure obtains worse results, but the lower number of seeds (18 seeds when there is one seed per depot per day) makes the customer assignment phase a lot faster. The time for this phase reduces from around 1200 seconds (20 minutes), to 260 seconds (4,3 minutes).

A disadvantage of the Period Scheduler clustering method with 18 seeds is that all vehicles go to one area on each day. One of the solution requirements is that the vehicle should go to multiple areas on one day, to cope with emergency orders.

Another disadvantage is that the workload is not spread perfectly across the depots and days, but this can be solved by improving the seed selection procedure, with a check if every depot has a seed for every day of the planning period. This improvement would make this seed selection procedure more interesting.

6 Conclusions and recommendations

This research is done by ORTEC's Oil, Gas, and Chemicals department. It has the objective to design a solution methodology to minimize distribution costs in the Inventory Routing Problem (IRP) for gas distribution, and to mitigate the seasonal peak in customer deliveries. We develop an algorithm to reach this objective, and we test the algorithm in computational experiments. In this chapter, we discuss the main conclusions, scientific contributions, recommendations, and ideas for future research.

6.1 Conclusions

The flexibility offered by Vendor Managed Inventory (VMI) in gas distribution can be fully utilized by integrating the inventory and routing decisions. In the current situation, the two decisions are separated by the two software suites offered by ORTEC: ORION (inventory) and SHORTREC (routing). We develop an algorithm based on the requirements in the gas industry to integrate the inventory and routing decisions.

The algorithm consists of three distinct phases, where the first phase selects all customers that can receive a relatively large delivery compared to the customer's capacity, the *relative delivery size*. The second phase creates clusters of customers around the customers that require a delivery in the planning period: the 'must-go' customers. The third phase assigns all customers to a delivery day in the planning period, with the objective to minimize total distance between customers. The algorithm considers the impact of a certain delivery on future planning periods, by evaluating the relative delivery size of customers that do not require a delivery in the planning period: the 'may-go' customers.

An equal workload balance improves the efficient use of the available resources. In the third phase, we also balance the workload on the short-term and the long-term, by adding a delivery volume constraint. This delivery volume constraint is connected to the vehicle capacity on a certain day. The maximum delivery volume for a planning period is calculated by averaging the forecasted weekly demand in the next n weeks. This moving average in demand flattens the demand curve and mitigates the seasonal peak. We test the algorithm with n equal to 13 and 26, where n equal to 26 flattens the curve more.

We compare the results of the algorithm plans with plans created by planners at a gas company. These tests illustrate that the algorithm delivers 21% more *volume per kilometre* than the actual plans, and decreases costs for gas distribution with at least 12% (the decrease in distribution costs due to fewer visits is not considered in this percentage). Additionally, all customers require fewer visits, which not only decreases distribution costs, but also improves customer satisfaction. The algorithm uses between 600 to 1.000 seconds to run, which is acceptable in a region with three depots, 3.000 customers, and a planning period of seven days. The four main benefits of the solution are:

- (1) Decrease in distribution costs, due to increased efficiency of route plans and fewer customer visits.
- (2) Increase in customer satisfaction, because there are fewer customer visits.
- (3) Improved workload balance in the planning period.
- (4) A mitigation of the seasonal peak in demand in the winter months.

Concluding, we provide a solution methodology that integrates the inventory and routing decisions to effectively utilize the business potential of VMI. This solution methodology improves the efficiency of gas distribution and mitigates the seasonal peak in demand.

6.2 Recommendations

Based on the conclusions from the previous paragraph, we state the long-term and short-term recommendations. The short-term recommendation is a short-term solution that is relatively easy to implement.

6.2.1 Long-term recommendations

- (1) The algorithm should be implemented in ORION and SHORTREC, where ORION is responsible for selecting the orders, the clustering, and the generation of schedules, and SHORTREC performs the assignment of the customers to delivery schedules. The single interaction between the two programs in this case decreases the complexity of the solution for the planners. Planners can work with the software suites in a similar way as they currently do. A plan should be created for at least a number of days, to get an equal workload. One can use the routes for the first day only, to get the effect of a rolling horizon, or use the routes of the complete planning period. The latter is efficient, because a planner can run the algorithm on a complete region for a planning period of seven days in 10 to 18 minutes.
- (2) ORION should calculate the demand of non-VMI customers as well. This is especially interesting if a peak in demand occurs for non-VMI customers. The demand of these customers can then be used to calculate a better average demand per week, and thus we improve the anticipation on demand peaks to come. To implement this idea, ORION should separate the functions of order generation and demand calculation, which are currently combined.

6.2.2 Short-term recommendation

The customer selection procedure proposed in this research can be used to select the set of possible ‘may-go’ customers in the planning process. In this way, the quality of the long-term planning is preserved and customer visits are kept to a minimum, since we only consider customers with a relative delivery size above α . In our computational experiments, we found that is difficult to find settings for α for every period of the year, due to peaks in demand and adjusted usage curves. The problem of setting α is avoided when we always select the 400 customers with the highest values for α , since 400 customers is the number of customers we found to yield good results in the algorithm.

6.3 Scientific contributions

- (1) The IRP with multiple depots is not described in the literature before. Several articles describe a solution with satellite facilities (Bard et al., 1998; Jaillet et al., 2002), but at these facilities vehicles can not start or finish a day, they can only reload. We have created an algorithm that balances the workload over multiple depots according to the available vehicle capacity at a depot.
- (2) The IRP with a seasonal peak in demand is described in the literature (Dror and Ball, 1987), but no author ever proposed solutions to mitigate this seasonal peak. Our algorithm uses a methodology that balances delivery volumes over a longer period, and thus flattens the demand curve effectively.
- (3) The algorithm works with a flexible clustering methodology, which is performed every time the algorithm is run. The literature only describes clustering methodologies for the IRP that are separated from the planning process and need to be re-run every time a change is made to customer data (usage curves, new customers, etc.). Therefore, the existing clustering methodologies are difficult to use in a real-life problem instance, where customer data are updated regularly. Flexible clustering provides a solution for real-life instances, since it is computationally fast.

- (4) This research proposes a solution for the problem of route failures. A route failure is a difference between the expected delivery volume in a trip and the actual delivery volume. A route failure results in either too little or too many customers in the pre-planned trips. The problem decreases utilization and this leads to a lower efficiency. Although Dror and Trudeau (1988) use the notification of route failures in their stochastic models, no author has described a practical solution to cope with these route failures. We propose balance customers with a flexible delivery volume, which are planned at the end of a trip, to cope with these route failures.
- (5) As mentioned earlier, the algorithm balances the workload according to vehicle capacity at a specific depot on a specific day in the planning period. This type of workload balancing in the IRP, based on the distribution of vehicle capacity in the planning period, is never described in the literature before.

6.4 Future research

- (1) We should test the algorithm in a feedback loop of multiple planning periods, where delivery information from one planning period is fed back into ORION, to calculate new due dates for following planning periods.
- (2) We propose balance customers with a flexible delivery volume to cope with route failures. We should study the balance customers in a stochastic environment, where route failures can be simulated, to draw conclusions about the actual effects.
- (3) We can improve computation times with the Period Scheduler seed selection procedure (Appendix A). The seed selection procedure should be improved to have an equal balance of seed points in the planning period. Additionally, the procedure selects a fixed region per depot and day, and this makes it difficult to plan emergency deliveries. We found that the number of emergency deliveries is not as high as expected ($\sim 0,8\%$ of all orders), so this may not pose a large problem.
- (4) We use flexible clustering in the algorithm; clusters are created based on the ‘must-go’ customers, the so-called seeds. The algorithm selects at least a certain minimum number of seed options per day in the planning period, by matching the latest delivery date of a customer with the current planning day we are selecting the minimum number of seed options for. We should adapt the current method, such that not only customers with a latest delivery date that matches exactly with the current planning day are considered, but also the customers that have a latest delivery date later than the current planning day.
- (5) To solve large ILPs, the current LP solver works with rounding, where the outcomes of the LP relaxation are rounded up. A better method, such as Lagrangian relaxation can improve, above all, the control of the short-term workload balance. A future research at ORTEC should study the use of Lagrangian relaxation to solve the ILPs.
- (6) We should study further integration of the two software suites ORION and SHORTREC. For instance, the swapping of orders between days after scheduling may be beneficial for the efficiency of trips. Currently, this is not possible, since ORION holds the required usage data, while SHORTREC schedules the routes.
- (7) We should study the use of forced orders for customers with a lower visit frequency than six times per year. We can force these orders in quiet times, such that these customers do not require a visit in the winter peak.
- (8) In evaluating the impact of a ‘may-go’ delivery on future planning periods, the ILP uses a forwarding cost function. We propose a tuning parameter for this cost function. For instance, in fall, the total delivery volume should be higher than in spring. Therefore, the costs for forwarding a customer should be lower in fall than in spring, such that fewer ‘may-go’ orders are forwarded to a future planning period, and the delivery volume is higher. This flattens the curve over the course of a year.

References

- Bard, J.F., Huang, L., Jaillet, P., and Dror, M. 1998. *A Decomposition Approach to the Inventory Routing Problem with Satellite Facilities*. Transportation Science 32 (2), 189-203.
- Bell, W.J., Dalberto, L.M., Fisher, M.L., Greenfield, A.J., Jaikumer, R., Kedia, P., Mack, R.G., and Prutzman, P.J. 1983. *Improving the Distribution of Industrial Gases with an On-line Computerized Routing and Scheduling Optimizer*. Interfaces 13, 4-23.
- Bertazzi, L., Paletta, G., and Speranza, M.G. 2002. *Determining Order-Up-To Level Policies in an Inventory Routing Problem*. Transportation Science 36 (1), 119-132.
- Blumenfeld, D.E., Burns, L.D., Daganzo, C.F., Frick, M.C., and Hall, R.W. 1987. *Reducing Logistics Costs at General Motors*. Interfaces 17 (1), 26-47.
- Brown, A.O., Ettl, M., Lin, G.Y., Petrakian, R., and Yao, D.D. 2002. *Inventory Allocation at a Semiconductor Company*. International Series in Operations Research and Management Science 42, 283-310.
- Campbell, A.M., Clarke, L., Kleywegt, A., and Savelsbergh, M.W.P. 1997. *The Inventory Routing Problem*. Fleet Management and Logistics, 95-113.
- Campbell, A.M., Clarke, L., and Savelsbergh, M.W.P. 2002. *Inventory Routing in Practice*. The Vehicle Routing Problem, SIAM monographs on discrete mathematics and applications, 309-330.
- Campbell, A.M., and Savelsbergh, M.W.P. 2004. *A Decomposition Approach for the Inventory-Routing Problem*. Transportation Science 38 (4), 488-502.
- Clarke, G., and Wright, J. 1964. *Scheduling of vehicles from a central depot to a number of delivery points*. Operations Research 12, 568-581.
- Dantzig, G.B., and Ramser J.H. 1959. *The truck dispatching problem*. Management Science 6, 80-91.
- Dror, M., and Ball, M. 1987. *Inventory/Routing: Reduction from an Annual to a Short-Period Problem*. Naval Research Logistics 34, 891-905.
- Dror, M., and Trudeau, P. 1988. *Inventory Routing: Operational Design*. Journal of Business Logistics 9 (2), 165-183.
- Federgruen, A., and Zipkin, P. 1984. *A Combined Vehicle and Inventory Allocation Problem*. Operations Research 32 (5), 1019-1037.
- Fisher, M.L. 1981. *The Lagrangian Relaxation Method for Solving Integer Programming Problems*. Management Science 27 (1), 1-18.
- Gardner Jr., E.S., Diaz-Saiz, J. 2002. *Seasonal Adjustment of Inventory Demand Series: A Case Study*. International Journal of Forecasting 18 (1), 117-123.

- Garey, M. R., and Johnson, D. S. 1979. *Computers and Intractability. A Guide to the Theory of NP-Completeness*. W. H Freeman and Company.
- Golden, B., Assad, A., and Dahl, R. 1984. *Analysis of a Large Scale Vehicle Routing Problem with an Inventory Component*. Large Scale Systems 7, 181-190.
- Hall, R.W. 1992. *A Note on Bounds for Direct Shipping Costs*. Management Science 38, 1212-1214.
- Hoendervoogt, A. (2006). *Period Scheduler*. MSc Thesis, ORTEC.
- Jaillet, P., Bard, J.F., Huang, L., and Dror, M. 2002. *Delivery Cost Approximations for Inventory Routing Problems in a Rolling Horizon Framework*. Transportation Science 36 (3), 292-300.
- Jung, J.H., and Mathur, K. 2007. *An Efficient Heuristic Algorithm for a Two-Echelon Joint Inventory and Routing Problem*. Transportation Science 41 (1), 55-73.
- Kleywegt, A.J., Nori, V.S., and Savelsbergh, M.W.P. 2004. *Dynamic Programming Approximations for a Stochastic Inventory Routing Problem*. Transportation Science 38, 42-70.
- ORTEC website. 2008. <http://www.ortec.com/solutions/vrd.aspx>. Accessed August 2008.
- Poot, A., Kant, G., and Wagelmans, A.P.M. 2002. *A Savings Based Method for Real-life Vehicle Routing Problems*. Journal of the Operational Research Society 53, 57-68.
- Song, J.H., and Savelsbergh, M.W.P. 2007. *Performance Measurement for Inventory Routing*. Transportation Science 41 (1), 44-54.
- Welch, T.H., Smith, J.G., Pix, J.P., and Reader, R.D. 1971. *Meeting Seasonal Peak Demands for Natural Gas*. Operational Research Quarterly 22, 93-106.
- Wieringa, R. and Heerkens, H. 2003. *Requirements Engineering as Problem Analysis: Methodology and Guidelines*. Working paper. University of Twente.
- Winston, W.L. 1994. *Operations Research; Applications and Algorithms*. Duxbury Press, Third Edition.
- Yugang, Y., Haoxun, C., and Feng, C. 2008. *A New Model and Hybrid Approach for Large Scale Inventory Routing Problems*. European Journal of Operational Research 189, 1022-1040.

Appendices

Appendix A: The Period Scheduler: Seed selection procedure	53
Appendix B: The Period Scheduler: ILP model	54
Appendix C: Extended literature summary	55
Appendix D: Forecasted and actual customer usage curves	65
Appendix E: Symbols used in the mathematical formulations.....	68
Appendix F: Workload calculation: Step 4.....	70
Appendix G: Workload calculation: Step 5.....	72
Appendix H: Workload calculation: Improving the balance.....	75
Appendix I: Alternative designs: Balance customers.....	76
Appendix J: Alternative designs: Different objective functions.....	80
Appendix K: Alternative designs: ILP with combination check	82
Appendix L: Settings used in the computational tests.....	83
Appendix M: Dataset overview.....	84
Appendix N: Comparison actual routes and SHORTREC routes.....	85
Appendix O: Computational results: The weekly comparison	86
Appendix P: Computational results: The working time balance.....	87
Appendix Q: Computational results: A different vehicle set-up.....	88
Appendix R: Computational results: α parameter	89
Appendix S: Computational results: Bandwidth parameter	90
Appendix T: Computational results: Distance type parameter	91
Appendix U: Computational results: Alternative designs.....	93

Appendix A: The Period Scheduler: Seed selection procedure

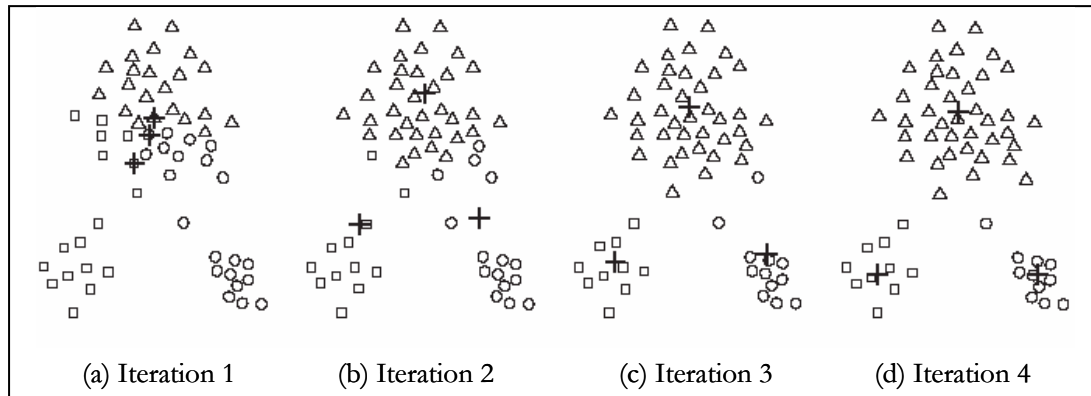


Figure A.1. Iterations to find the seed points for the Period Scheduler.

Figure A.1 illustrates an example of the seed selection procedure. The crosses indicate the seed points, and the triangles, circles, and rectangles indicate the customers. The procedure selects three random seed points and assigns every customer to one seed point, with the objective to minimize the total sum of the distances between the customers and the seed point. A different shape indicates a different seed assignment (Figure A.1.a). Based on this information, new seed points are calculated by calculating the centre of gravity for each of the three regions, and setting this centre as the new seed point. The customer assignment is repeated with the new seed points (Figures A.1.b and A.1.c), and again distance is minimized. The last step is repeated until the seed points stop changing (Figure A.1.d). The Period Scheduler prevents large differences in the total delivery volume per region by basing the assignment of a customer to a seed point not only on distance, but also on delivery volume.

Appendix B: The Period Scheduler: ILP model

$$\min \sum_{i=1}^n \sum_{j \in S_i} d_{ij} X_{ij} \quad (\text{B.1})$$

s.t.

$$\sum_{j \in S_i} X_{ij} = 1 \quad \text{for customer } i = 1, \dots, n \quad (\text{B.2})$$

$$\sum_{i=1}^n \sum_{j \in S_i} q_{ijk} X_{ij} \leq w_{\max k} \quad \text{for day } k = 1, \dots, K \quad (\text{B.3})$$

$$X_{ij} \in \{0,1\} \quad \text{for customer } i = 1, \dots, n \text{ and } j \in S_i \quad (\text{B.4})$$

Variables

X_{ij} binary variable that equals 1 if delivery schedule j is assigned to customer i

Data

d_{ij} distance contribution (i.e. costs) to assign delivery schedule j to customer i

$w_{\max k}$ maximum workload on day k in terms of size m

q_{ijk} quantity delivered on day k , when schedule j is assigned to customer i

S_i set of allowable delivery schedules for customer i

n number of customers

J number of different delivery schedules

K number of days

Figure B.1. The ILP model in the Period Scheduler (Hoendervoogt, 2006)

The ILP model in the Period Scheduler has the objective to minimize distribution costs in Equation B.1. Constraint B.2 guarantees that we have a schedule for each customer. Constraint B.3 balances the workload per day by setting the total sum of the delivery volume lower than or equal to the maximum workload. The binary property of the variable X_{ij} is set in Constraint B.4.

Appendix C: Extended literature summary

This appendix summarizes the used literature in this thesis. Paragraph C.1 discusses key principles, which form the basis for finding a solution methodology for any type of IRP. Additionally, we discuss nine solution methodologies in Paragraph C.2. Paragraph C.3 discusses performance measurement for the IRP, and Paragraph C.4 discusses the seasonal peak problem in the literature.

C.1 Key principles in solving the IRP

A solution methodology for the IRP is not straightforward. Four important articles describe the problem and propose key principles, or guidelines, for designing a solution methodology. Dror and Ball (1987) start with analyzing a single customer case. In order to solve the problem, they define cost functions for orders that fall either within or outside the planning period. The two types of orders they describe, correspond exactly with the mentioned ‘must-go’ and ‘may-go’ orders used in this thesis. Callego and Simchi-Levi (1990) study the strategy to only consider routes that exist of one customer: a direct-shipping strategy. Campbell et al. (1997) discuss the single customer case as well. The authors also give two strong guidelines that are very important in designing a solution for the IRP. Yugang et al. (2008) describe an interesting measure for defining travel distance, which is the main aspect of transportation costs in their approach. The articles are ordered by publication year.

C.1.1 Dror and Ball (1987)

Dror and Ball (1987) analyze the IRP for a supplier in heating oil. The authors observe a seasonal peak in customer usage and incorporate it into the customer usage functions. The authors state that the objective of the IRP is to minimize the long-term distribution costs, while making sure no customer runs out of stock. Two major practical problems arise when trying to reach this objective for a long period:

- (1) The validity of a formulation over a long period of time is questionable because parameter values are uncertain over such a long period (new customers, etc.).
- (2) The number of constraints and variables is prohibitively high.

Therefore, the procedure for reducing the long-term optimization problem and selecting the set of customers that is actually replenished in the short-term is of major importance in developing a solution for the IRP. According to the authors, the key element in reducing the long-term optimization problem to a short-term problem is to include penalty costs within the short-term model that reflect the long-term effect of a decision.

The authors consider the single customer problem. In this case, the optimal policy would be to refill the customer’s tank exactly at the time it would reach a zero level. The optimal t^* can be calculated by dividing the inventory level at the beginning of the considered period I^0 by the daily usage μ :

$$t^* = \frac{I^0}{\mu} \quad (C.1)$$

The authors model the cost increase when a customer is not replenished at t^* , but at t ($t < t^*$). They define two groups of customers, one group where the optimal delivery falls within the considered time period of m days: ‘must-go’ ($m \cdot \mu - I^0 < 0$), and a second group where the optimal delivery falls outside the considered time period: ‘may-go’ ($m \cdot \mu - I^0 \geq 0$).

They define the cost for a ‘must-go’ order as the optimal volume (I^0) minus the delivery we can do if we pull the order forward (μt). The ‘missed’ delivery volume is then the cost for pulling an order forward. They divide this volume by the customer’s tank capacity, C , to get the relative size of the missed delivery as a percentage of capacity, which represents the long-term effect of a decision to deliver on the short-term. Equation C.2 illustrates the cost for not delivering the optimal volume, c_t .

$$c_t = \frac{I^0 - \mu t}{C} b \quad (C.2)$$

Where b is the cost for a delivery. Equation C.2 shows that the authors minimize c_t by maximizing the amount delivered and choosing a delivery day t as close to t^* as possible.

The ‘may-go’ orders do not receive a penalty, but an incentive to force these orders in the considered time period. This is because their model has no other incentive to plan the ‘may-go’ orders in the considered time period. The authors use Equation C.3 to calculate the cost g_t for planning a ‘may-go’ order in this period, under the assumption that the optimal policy is used afterwards. This cost g_t is actually a reward, since it has a negative sign in the overall cost function, thus g_t is a *decrease* in total cost. The equation to calculate g_t is based on c_p , and the smaller c_p , the larger the incentive to plan a ‘may-go’ order. Equation C.2 defines c_p , and we can see that c_t is small when I^0 is small and t is large. *Therefore, it is more attractive to pull forward a ‘may-go’ order that is close to its optimal delivery point (large t).* It is also more attractive to pull an order forward from a customer with a small initial inventory (small I^0). The model assigns as many ‘may-go’ orders as possible, since g_t is a *decrease* in total cost, and c_t is smaller than b for all ‘may-go’ orders.

$$g_t = b - c_t \quad (C.3)$$

C.1.2 Callego and Simchi-Levi (1990)

Callego and Simchi-Levi (1990) investigate the long-term effectiveness of direct shipping (separate loads to each customer). They conclude that direct shipping is at least 94% effective over all inventory routing strategies when the minimal economic lot size is at least 71% of truck capacity. The effectiveness deteriorates as the economic lot size gets smaller. Hall (1992) points out that direct shipping only performs well when fixed transportation costs are negligible.

C.1.3 Campbell, Clarke, Kleywegt, and Savelsbergh (1997)

Similar to Dror and Ball (1987), the authors state that in the single customer problem, it is optimal to refill the customer’s tank precisely at the time it becomes empty. The cost v_t for replenishing a single customer for a planning period of time T is calculated as follows:

$$v_t = \max(0, \left\lceil \frac{Tu - I}{\min(C, Q)} \right\rceil) c. \quad (C.4)$$

Where C is the customer capacity and Q is the vehicle capacity. The delivery costs are given by c , and the usage is defined by u . The initial inventory is given by I .

For the multiple customer problem, the decision who needs to be visited, and how much should be delivered, should be guided by the following assumptions:

- (1) Always try to maximize the quantity delivered.
- (2) Always try to send out vehicles with a full load.

These assumptions are straightforward, since maximizing the quantity delivered will yield a better performance on the long run, as shown in the single customer problem in Dror and Ball (1987) and Campbell et al. (1997). Since we incur fixed costs when we send out a vehicle, its effectiveness is maximized by only sending out the vehicle with a full load.

C.1.4 Yugang, Haoxun, and Feng (2008)

The authors point out that transportation costs include the fixed usage costs that are related to vehicle insurance, depreciation and rewards for drivers, but also the variable transportation costs depending on the travel distance and the delivery quantity. The authors make use of the *triangle inequality*, where, in a set of two customers and a depot, a trip through the depot and the two customers is always larger than, or at least as large as, a trip between the depot and one of the two customers. This is illustrated in Figure C.1, a trip between the Depot and Customer 1 will be: $2 * 100 = 200$, and a trip including Customer 2 will be: $100 + 30 + 90 = 220$. This means we increase distance by adding customers that are far from Customer 1, but that the increase is only to be measured as distance of the trip with one depot and two customers minus the distance of the trip with one depot and one customer. In Figure C.1, the additional distance when Customer 2 is added to the trip would be: $90 + 30 - 100 = 20$, or: $220 - 200 = 20$.




Figure C.1. Trips can be seen as triangles.

C.2 Solution methodology for the IRP

Bell et al. (1983), Federgruen and Zipkin (1984), Golden et al. (1984), and Blumenfeld et al. (1987) are among the first authors to describe a solution methodology for the IRP. Since then, researchers have tried to define heuristics to solve the IRP for different problem instances. The main differences in these approaches are:

- (1) A deterministic versus a stochastic approach, where customer usage is assumed to be either deterministic or stochastic.
- (2) A decomposition versus an integrated approach, where a decomposition approach tries to find a solution for the IRP in a phased approach, where the first phase will result in a time and quantity for a delivery, and the second phase solves the resulting VRPs. The integrated approach deals with the decision when and how much to deliver, and the VRPs, at the same time.
- (3) Different decision models are used in the literature, where most authors use an ILP to minimize cost or maximize revenue; other authors solve the IRP with Dynamic Programming (DP), or with heuristics.

Additionally, many differences in problem aspects arise, such as length of the planning period, number of products, inventory holding costs, safety stock levels, maximum stock levels, fleet aspects, multiple reload facilities, size of the customer instance, and workload balancing. Table C.1 illustrates an overview of the articles discussed in this paragraph. We discuss the authors in the order of year of publication.

Author(s)	Planning period	Industry	Customer usage	Number of customers	Satellite facilities	Decomposition or Integrated	Decision model
Golden, Assad, and Dahl (1984)	1 day	Liquid propane	Deterministic	3000	No	Decomposition	-
Dror and Ball (1987)	1 week	Heating oil	Stochastic	>1000	No	Decomposition	ILP
Dror and Trudeau (1988)	1 week	Heating oil	Stochastic	2077	No	Decomposition	ILP
Bard, Huang, Jaillet, and Dror (1998)	1 week	Liquid propane	Deterministic	500	Yes	Decomposition	ILP
Jaillet, Bard, Huang, and Dror (2002)	1 week	Liquid propane	Deterministic	500	Yes	Decomposition	ILP
Bertazzi, Paletta, and Speranza (2002)	30 days	Echelon network	Deterministic	50	No	Integrated	Heuristic
Campbell and Savelsbergh (2004)	3 days	Liquid propane	Deterministic	150	No	Decomposition	ILP
Kleywegt, Nori, and Savelsbergh (2004)	1 day	Liquid propane	Stochastic	20	No	Integrated	DP
Jung and Mathur (2007)	-	Echelon network	Deterministic	1000	No	Decomposition	Nonlinear IP

Table C.1. An overview of the solution methodologies discussed in Paragraph C.2.

C.2.1 *Golden, Assad, and Dahl (1984)*

Golden et al. (1984) investigate the deterministic IRP at a large propane distribution company. They consider a district with 3.000 customers and have little data to base forecasts and calculations on. They propose an integrated solution methodology where the *relative delivery size* for all customers is calculated by dividing the possible delivery to the customer by its capacity, the left term in Equation C.5. All customers that can receive a minimum relative delivery size in the planning horizon are considered in the current planning horizon. This minimum relative delivery size can be set with α in Equation C.5.

$$\frac{\text{PossibleDeliveryVolume}}{\text{Capacity}} \geq \alpha \quad (\text{C.5})$$

If Equation C.5 holds for a certain customer and certain α , the customer is selected for the potential customer set in the current planning horizon. The authors study the sensitivity of the setting of α . When α is close to 1, orders are only accepted for customers that have no gas left, meaning the number of visits is lower, and thus transportation costs are lower, but the number of stock outs will rise, resulting in an overall increase in distribution costs. The authors claim that stock outs are expensive since there are two cost factors: (1) the high cost of an emergency delivery, and (2) the extra cost for an unhappy customer. When α is close to 0,5, the final solution will be more expensive, since we put orders in the potential customer set that can only receive 50% of their optimal delivery volume.

Golden et al. (1984) use the following solution methodology:

- STEP 1:** Select customers for the potential customer set by setting α .
- STEP 2:** Calculate distance and profit matrix for all potential customers.
- STEP 3:** Select customers in a trip by adding a customer to a TSP tour with a maximum length: a TSP with Time Constraint (TSPTC).
- STEP 4:** Construct trips with Clarke-Wright algorithm².
- STEP 5:** Assign trips to trucks.
- STEP 6:** Check feasibility, if not feasible, add overtime and/or increase the maximum length in Step 3.

The authors describe two dimensions along which the optimization in an IRP must proceed. The authors state that the IRP optimization is done on a *spatial* dimension and a *temporal* dimension. The spatial dimension requires the minimization of distance travelled. The temporal dimension involves the timing of deliveries so that both early and late deliveries are discouraged, since these deliveries are inefficient. For the temporal dimension, they use the relative delivery size.

C.2.2 Dror and Ball (1987)

Most of the article by Dror and Ball (1987) is discussed in the Paragraph C.1.1. Their algorithm solves the stochastic IRP in three steps:

- STEP 1:** Assign customers to days solving an assignment problem with an LP.
- STEP 2:** Solve a VRP for each day using a modified Clarke-Wright algorithm.
- STEP 3:** Improve the solution by inter-route and inter-day customer exchanges.

C.2.3 Dror and Trudeau (1988)

Dror and Trudeau (1988) study the stochastic IRP for the distribution of propane. The authors claim that the actual optimal replenishment should occur a certain probabilistic time before the customer becomes empty, since the cost curve rises steeply in these last moments due to the risks of a stock out. They also consider route failures in their probabilistic analysis. A route failure is a failure, where the vehicle's route has either too little or too many customers. Since the propane industry works with a 'top-up' policy, in which a customer is filled up to its maximum capacity, the actual delivery to a customer is only known after the delivery is made. Therefore, a delivery might be larger or smaller than expected, and a route failure can occur, since the vehicle is out of stock before it reaches the last client in the trip or still has stock after visiting the last client.

C.2.4 Bard, Huang, Jaillet, and Dror (1998)

Bard et al. (1998) propose a heuristic to solve the deterministic IRP with satellite facilities in the propane industry. They consider a planning horizon of two weeks and they select all customers that have to be visited in this period. Additionally, they balance the daily delivery volume and schedule all routes in the m homogeneous vehicles for the first week only. The planning horizon is then shifted with a week and the complete process is repeated. The assignment problem is solved with a Mixed Integer Linear Program problem (MILP) with two objectives: (1) minimize the distance travelled, and (2) minimize the added costs that arise because the customer is not delivered at the optimal

² Since the routing step is out of scope for this thesis, we will not explain the Clarke-Wright algorithm, the reader is referred to Clarke and Wright (1964).

point in time, the *incremental costs*. An MILP is an LP, where a subset of decision variables has an integer constraint. Rather than combining the two measures into a single value, the authors optimize in two steps. Although it may be possible to translate distance into cost, the results may not be directly additive with the incremental costs, because the incremental costs may have an intangible aspect corresponding to quality of service and customer retention.

Their heuristic solves the IRP in five steps:

- STEP 1:** Identify all customers whose optimal replenishment day falls within the planning horizon of two weeks.
- STEP 2:** Assign each customer from Step 1 to a day by solving a balanced assignment problem with an MILP, with the objective of minimizing incremental costs.
- STEP 3:** For each day in the planning horizon, find a good solution by solving a VRP with Satellite Facilities (VRPSF), with the objective to minimize total distance.
- STEP 4:** Improve the solution by swapping customers between routes.
- STEP 5:** Improve the solution by examining the trade-off between incremental costs and route lengths, by swapping customers between different days of the plan.

A VRPSF is a VRP with satellite facilities. A satellite facility differs from a depot in that a vehicle can only reload at a satellite facility, and it can not start or finish there. The heuristics tested in Step 3 – 5 are the Clarke-Wright algorithm, a GRASP algorithm, and a revised Sweep Algorithm. The Clarke-Wright algorithm performs best in a VRPSF, but all algorithms have acceptable computation times.

C.2.5 *Jaillet, Bard, Huang, and Dror (2002)*

Jaillet et al. (2002) use an identical approach as Bard et al. (1998). The objective is to minimize cost on the long run, by determining the cost for delivering the order earlier than the optimal delivery interval for every customer in a planning period of n days. The authors determine the expected cost in case the delivery interval is not the optimal interval, but a shorter or longer period.

C.2.6 *Bertazzi, Paletta, and Speranza (2002)*

Bertazzi et al. (2002) consider a deterministic IRP in a two-echelon network with one depot, one vehicle, and multiple retailers. They consider holding costs, maximum inventory levels, and they accept no stock outs. They propose a heuristic to solve the IRP with an integrated approach. They generate a list of all customers ranked in the non-decreasing order of the average number of time units needed to consume their inventory, which is calculated with Equation C.6.

$$\frac{\text{MaximumStock} - \text{SafetyStock}}{\text{DailyUsage}} \quad (\text{C.6})$$

With the list, they generate a starting solution, by adding customers to a vehicle trip at time t , where the interval between deliveries must be equal to the number of time units the customer needs to consume their inventory. After they have filled all vehicles, and planned all customers, they improve the solution by removing customers from trips, and inserting the customer in a new trip at time t' , where t' is calculated with the knowledge from the initial solution.

C.2.7 Campbell and Savelsbergh (2004)

Campbell and Savelsbergh (2004) propose a method to solve the deterministic IRP in the commercial gas industry. They propose two separate phases. In Phase 1, deliveries are assigned to days over the k -days horizon. In Phase 2, deliveries are assigned to days and minutes, over the j -days horizon, where $j \leq k$, and trips are scheduled. The decisions made in Phase 1 are not engraved in stone, but are used as guidelines.

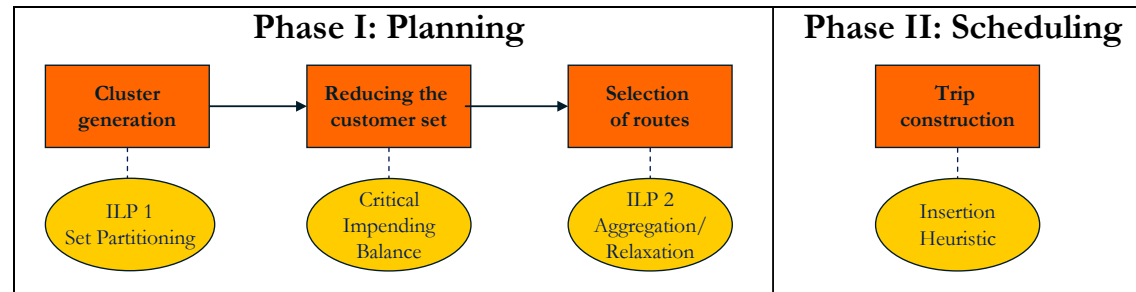


Figure C.2. The solution methodology of Campbell and Savelsbergh (2004)

Phase I: Assign deliveries to days over the k -day horizon

In Phase I, they create routes, containing customers that are served together, and they choose the best routes with an ILP. A route is different from a trip, in that a route does not prescribe the order in which the customers should be served. The objective is to minimize total distribution costs, and the cost of a route c_r is considered to be the distance of the optimal TSP-tour through all customers in the route. All deliveries in one route on one day, can not exceed the homogeneous truck capacity, and the total work that can be done is restricted by the total working time available.

Clustering

A problem in this approach is the large number of available routes; this was also noted by Bell et al. (1983). To limit the number of possible routes, groups of customers are created, which are called clusters. Clusters are groups of customers that can be served cost effectively by a single vehicle for a long period of time. Clusters are constructed after a change to the customer set, for instance a change of customer usage patterns. The following approach is used to identify a set of disjoint clusters covering all customers:

- (1) Generate a large set of possible clusters.
- (2) Estimate the costs of serving each cluster.
- (3) Solve a set-partitioning problem to select clusters.

Cluster costs are estimated by calculating the distribution costs for serving the customer in a cluster for a certain period. An ILP is used to generate the costs for a cluster, and the reader is referred to Campbell and Savelsbergh (2004) for this ILP. The minimal TSP tour through all customers in a cluster should be below a certain time threshold, but cluster costs do not only depend on distance, but also on compatible inventory capacities and usage rates of the customers. For example, five customers that need a full truckload per day will not be combined into one cluster.

Reducing the customer set

Additionally, the authors reduce the customer to reduce the problem size. Campbell and Savelsbergh (2004) identify several groups that are included. The first group are *critical* customers. Critical customers have a large impact on the efficiency of the schedule and include those customers that have high demand or are very distant from any other

customers or the depot. The second group are the *impending* customers, which require a delivery in the next several days. The third group are the *balance* customers to improve workload balance. Balance customers are customers that do not need a delivery imminently, but are near and in the same cluster as those customers that are critical or impending. The b nearest neighbours for every critical or impending customer within the cluster are included, that can receive a delivery of a minimum size in the next few days. Phase I has a higher probability for finding good routes when these balance customers are included.

Phase II: Scheduling of routes

The routes that are selected in Phase I are then scheduled in Phase II. This is done for several days with an insertion heuristic. The Phase I results are considered as an advice in Phase II, to keep flexibility in scheduling. Since we do not consider the scheduling in our thesis, we do not discuss Phase II in detail. The reader is referred to Campbell and Savelsbergh (2004).

C.2.8 Kleywegt, Nori, and Savelsbergh (2004)

Kleywegt, Nori, and Savelsbergh consider a stochastic IRP with a single vendor, multiple customers, multiple vehicles, and a homogeneous fleet. An optimal value function is defined to find a solution to the IRP. In their formulation, the authors include a reward per litre delivered, costs for traversing a certain arc in the network, inventory holding costs, and stock out penalties.

To approximate the optimal value function, they decompose the IRP into sub problems, each designed to have two properties: (1) it provides an accurate representation of a portion of the overall problem, and (2) it is relatively easy to solve. The sub problems are problems that need to be solved for subsets of customers, and the combined sub problems give an accurate representation of the total problem. A customer can be in one or more subsets. They use Dynamic Programming (DP) to generate policies that are feasible and a good approach to the optimum. For each policy, they consider all feasible decisions which contain information about (1) when to deliver to a customer, (2) how much to deliver to a customer, and (3) how to combine customers in vehicle routes. For more information on the application of DP in the IRP, see Kleywegt et al. (2004).

To proof their findings, they simulate a problem instance with 50 customers and they consider subsets of one, two, and three customers only. The authors state that this is a good assumption in the petroleum and air industry, since routes usually contain one customer, and not more than three customers. Since the number of options rise, the solutions improve when the authors consider a higher number of customers per route. Table C.2 illustrates that a negative effect of a high number of customers per subset is a strong increase in computation times. The authors note that this is an approach where the complete customer set is considered, one could easily reduce the number of two or three customer routes by eliminating all subsets of two or three customer routes that can not be visited together.

Number of customers per route	Computation time
1 customer	0,5 seconds
2 customer	318,5 seconds
3 customer	668.360 seconds (7,73 days)

Table C.2. An overview of the computation times for solving all sub problems where there is a possibility to include one, two, and three customers in one sub problem.

C.2.9 Jung and Mathur (2007)

Jung and Mathur study a two-echelon deterministic IRP with one warehouse and N retailers. They consider the problem with holding costs at the retailer and at the warehouse, and homogeneous vehicles to visit the retailers. To solve the problem, they cluster the retailers in such a way, that all retailers in one cluster can be visited together for a longer period of time by a single vehicle. The clusters only have to be made when there is a change in the customer set. They use stationary policies, or fixed visit schemes, in which a cluster's reorder interval is exactly 0,5 or 2 times the reorder interval of the warehouse. According to the authors, this makes the problem tractable. The fixed reorder intervals form the basis of solving the complete problem, since solving the exact reorder intervals for every customer and for the warehouse is their objective.

C.3 Performance measurement for the IRP

An evaluation of a solution methodology for the IRP is not obvious. It is a difficult, NP-hard problem, and because of its complexity, we have no optimal solution to compare our heuristic with. In most cases, a lower bound can be calculated for an NP-hard problem, which is a solution to a problem that is a more simple derivation of the actual problem. A lower bound for the IRP can be calculated, but this is a complicated task and an increase in complexity of the IRP will yield a weaker lower bound. The weaker a lower bound, the less useful it becomes. Song and Savelsbergh (2007) discuss the lower bound for the IRP and the performance measures one can use in evaluating a solution methodology for the IRP.

C.3.1 Song and Savelsbergh (2007)

A popular performance measure for the IRP is the *volume per mile* measure. Since the volume that has to be delivered in the IRP is given by customer usage, a company tries to minimize the distribution cost to deliver this volume. An important driver for distribution cost is distance travelled, and therefore maximizing volume per mile is an important measure. Song and Savelsbergh (2007) state that this practical performance measure is very effective in measuring relative performance, but inadequate in measuring absolute performance, because geography of customer locations and customer usage patterns vary across problem instances. Despite the inadequacy in measuring absolute performance, the measure is valuable for monitoring the performance in a single region over time. If the region has a stable customer set, and customer usage patterns do not fluctuate much, then an increase in volume per mile indicates that distribution planning is improving.

The authors propose an upper bound for the IRP is weak in a problem instance when customer capacity is substantially smaller than vehicle capacity. For this reason, a more complex upper bound is developed, which is useful when the typical number of stops on a delivery route is small. When the typical number of stops on a delivery route is large, the computational requirements become prohibitive. The reader is referred to Song and Savelsbergh (2007) for this complex upper bound.

C.4 Seasonal peak problem

Where the IRP is well represented in the literature, the seasonal peak problem in distribution is a problem that has never been addressed before. Articles about forecasting inventory demand for a US auto parts distributor are available (Gardner and Diaz-Saiz, 2002), but the distribution strategy that evolves from a forecast has never been described. An article from Welch, Smith, Pix, and Reader (1971) studies the seasonal peak problem in the production of gas, and some of their concepts are applicable to the seasonal peak in the distribution of gas.

C.4.1 *Welch, Smith, Pix, and Reader (1971)*

The authors propose five solutions to cope with the seasonal peak problem in production. We transfer these solutions to distribution in the last section of this paragraph.

- (1) Stocking of goods prior to peak period.
- (2) Fluctuating production rates through the year, with consequent idle capacity at off-peak times.
- (3) Failing to meet some of the demands at peak times.
- (4) Producing a mix of products, such that the demand patterns of some are complementary to others.
- (5) Seasonal pricing or seasonal contracts to reduce demand at peak times.

Option (1) can solve the seasonal peak problem for distribution as well. When stock at the customer is sufficient to balance the usage in the peak period, no delivery has to be made in the peak period. Option (2) is currently a method to cope with this problem; resources are purchased based on the winter period, with idle capacity in summer. Option (3) is something we would like to prevent from happening. Option (4) is a production opportunity, but might be feasible in gas distribution as well. If the idle capacity in summer would be used to deliver fuel products (oil products) the deliveries might be balanced better throughout the year. Since this is outside the scope of this thesis, we will not consider this option further. Option (5) is also outside the scope of this thesis.

Appendix D: Forecasted and actual customer usage curves

This appendix illustrates the forecasted usage and actual deliveries for different groups of customers. The theoretical visit frequency (VF) for a year is calculated by dividing the total usage for a customer in a year by the maximum capacity minus the safety stock.




Figure D.1. Forecasted weekly usage and actual weekly delivery volumes of gas in the NE-region for all customers with a visit frequency of one visit per year.




Figure D.2. Forecasted weekly usage and actual weekly delivery volumes of gas in the NE-region for all customers with a visit frequency of two visits per year.




Figure D.3. Forecasted weekly usage and actual weekly delivery volumes of gas in the NE-region for all customers with a visit frequency of three visits per year.




Figure D.4. Forecasted weekly usage and actual weekly delivery volumes of gas in the NE-region for all customers with a visit frequency of four visits per year.




Figure D.5. Forecasted weekly usage and actual weekly delivery volumes of gas in the NE-region for all customers with a visit frequency of five visits per year.




Figure D.6. Forecasted weekly usage and actual weekly delivery volumes of gas in the NE-region for all customers with a visit frequency of six visits per year.




Figure D.7. Actual weekly delivery volumes of gas in the NE-region for all customers with no usage forecasted (non-VMI customers (grain farms, regular customers) and new VMI customers). The grain farm peak, due to increased drying of the harvested grain, is clearly visible around week 37.




Figure D.8. Forecasted weekly usage and actual weekly delivery volumes of gas in the NE-region for all customers with the seasonal profile 'Holiday Park'.

Appendix E: Symbols used in mathematical formulations

Variables used in the algorithm

X_{sep}	binary variable that equals 1 if delivery schedule s is assigned to customer c , and depot p
$X_{sck^{tp}}$	binary variable that equals 1 if delivery schedule s is assigned to customer c , and seed delivery k^{tp} , 0 otherwise
X_c	binary variable that equals 1 if customer c is excluded from the current planning horizon, 0 otherwise

Data used in the algorithm

y_{ts}	relative delivery size in percentage on day t in schedule s
\tilde{y}_c	relative delivery size in percentage if customer c is forwarded, the relative delivery size at the end of the next planning period
d_{cp}	distance contribution to assign customer c to depot p
$d_{ck^{tp}}$	distance contribution to assign customer c to seed delivery k^{tp}
\tilde{d}_c	is the distance from customer c to its nearest seed delivery
k^{tp}	seed delivery k that is assigned to day t , and depot p
$w_{\max tp}$	maximum number of seed deliveries on day t , depot p
$w_{\min tp}$	minimum number of seed deliveries on day t , depot p
$v_{\max k^{tp}}$	maximum workload for seed delivery k^{tp} on day t
$v_{\min k^{tp}}$	minimum workload for seed delivery k^{tp} on day t
m_{st}	1 if there is a delivery on day t in schedule s , 0 otherwise
r_{st}	integer variable for the delivery volume on day t in schedule s
S_c	set of allowable delivery schedules for customer c
$R_{sk^{tp}}$	volume for the deliveries after the first delivery in schedule s that is assigned to seed k^{tp}
SC	set of seed customers
SCD	set of seed customer deliveries
CC	set of clustered customers
CC_{mustgo}	set of clustered customers that are ‘must-go’ customers
D	set of depots
T	number of days in the planning horizon

Data used in workload calculation

VC_{tp}	available vehicle capacity at depot d on day t
F_c	number of deliveries in the planning period for customer c
β_{seed}	bandwidth used for seed customer workload calculation
β	bandwidth used for clusterer customer workload calculation
D_t^{total}	total number of depots opened on day t
S_c^{total}	total number of delivery schedules available to customer c
$w_{exp tp}$	total expected number of seed deliveries on day t and depot p
$v_{exp tp}$	total expected delivery volume on day t and depot p
Ω_{max}	total maximum delivery volume for a planning period
Ω_{mustgo}	total expected delivery volume for all ‘must-go’ customers
$\Omega_{forecasted}$	total expected average demand over the next n planning periods
Ω_{seed}	total expected delivery volume for all seed customers
$\Omega_{clustered}$	total expected delivery volume for all clustered customers

Data used in the alternative designs of the algorithm

BC	set of balance customers
λ, μ, κ	objective function setting parameters
$\varsigma_{k^p c}$	is 1 if the customer c can be assigned to seed delivery k^p , 0 otherwise

Appendix F: Workload calculation: Step 4

The maximum workload calculation for Step 4, the assignment of seed customers to schedules, is explained in Paragraph F.1, and the minimum workload calculation is explained in Paragraph F.2.

F.1 Maximum workload

The maximum workload per depot and day can be calculated with Equation F.1, which is rounded up.

$$w'_{\max tp} = \left\lceil \left(\frac{VC_{tp}}{\sum_{j \in D} \sum_{i=1}^T VC_{ij}} \right) \cdot \sum_{c \in SC} F_c \right\rceil \quad (F.1)$$

Where t indicates the day in the planning period, and p indicates the depot. VC_{tp} is the vehicle capacity on day t for depot p . SC is the set of seed customers, and D is the set of depots. T is the number of days in the planning period, and F_c is the number of deliveries for customer c in the planning period.

The maximum workload calculated with Equation F.1 can still lead to an infeasible ILP. A problem arises when there are more seed customers that need a delivery on Monday than the volume calculated with Equation F.1 for Monday. More specifically, the problem can even occur when n seed deliveries have to take place on Monday or Tuesday, but we have less than $0,5n$ maximum workload calculated for Monday, as well as Tuesday. Therefore, we calculate the minimum for the maximum number of seed deliveries by assuming that for a single customer every allowed schedule has an equal probability of being chosen. If we multiply this probability with a 1 if there is a delivery on the specific day, and a 0 if there is no delivery, we get the expected number of deliveries on a specific day for a single customer. Summing all expected numbers of deliveries over all customers results in Equation F.2, where $w_{\exp tp}$ is calculated:

$$w_{\exp tp} = \left(\frac{1}{D_t^{\text{total}}} \right) \sum_{c \in SC} \left(\frac{\sum_{s \in S_c} m_{st}}{S_c^{\text{total}}} \right) \quad (F.2)$$

Where m_{st} is 1 if there is a delivery on day t in schedule s , and 0 if there is no delivery on day t in schedule s . S_c^{total} is the number of schedules that are allowed for customer c . D_t^{total} is the number of opened depots on day t . The result of Equation F.2 is a number that equals the total expected number of deliveries for day t , and depot p .

The maximum number of seed deliveries per depot, day, and vehicle type set is the maximum of Equations F.1 and F.2:

$$w_{\max tp} = \max(w'_{\max tp}, \lceil w_{\exp tp} \rceil) \quad (F.3)$$

F.2 Minimum workload

The minimum number of seed deliveries is calculated with Equation F.4, which is a function that is similar to Equation F.1, the function to calculate the maximum workload with. The difference is that the equation is rounded down.

$$w'_{\min tp} = \left\lfloor (1 - \beta_{seed}) \cdot \left(\frac{VC_{tp}}{\sum_{j \in D} \sum_{i=1}^T VC_{ij}} \right) \cdot \sum_{c \in SC} F_c \right\rfloor \quad (F.4)$$

When the expected number of seed deliveries, equal to the value of a rounded down Equation F.2 is below the outcome of Equation F.4 for a day and a depot, the minimum number of seeds for this day and depot is set to the value of Equation F.2 rounded down. This is done to ensure the ILP is feasible, and the minimum workload constraint can be met. Equation F.5 illustrates this.

$$w_{\min tp} = \min(w'_{\min tp}, \lfloor w_{\exp tp} \rfloor) \quad (F.5)$$

F.3 Improving the balance of the workloads

When the minimum (or maximum) workload for a certain day and depot is calculated with Equation F.2, and not with Equation F.4 (or F.1), the total minimum (or maximum) workload does not equal the summation of all $w'_{\min tp}$ (or $w'_{\max tp}$) over all days and depots anymore. We have to correct all the workload calculations for the other days and depots, such that the total workload is equal to this summation again. This problem is also apparent in the assignment of clustered customers to delivery schedules, and Appendix H explains the method that is used for this procedure. It is best to read Appendix H after understanding the assignment of clustered customers in Step 5.

Appendix G: Workload calculation: Step 5

The workload is balanced evenly over the depots and the days in the planning period, depending on the available vehicle capacity at a depot on a specific day in the planning period. To obtain this, we calculate a maximum and a minimum workload.

G.1 Maximum workload

Equation G.1 illustrates the balancing of the volume of clustered customers. We calculate the workload for every seed customer delivery, and it is based on the day of the delivery and the depot assignment of the seed customer.

$$v'_{\max tp} = (1 + \beta) \cdot \left(\frac{VC_{tp}}{\sum_{j \in D} \sum_{i=1}^T VC_{ij}} \right) \cdot \Omega_{\max} \quad (G.1)$$

Where t indicates the day in the planning period, and p indicates the depot. β is a bandwidth to ensure the problem is feasible; it is also used in the minimum workload. VC_{tp} is the vehicle capacity on day t for depot p , and D is the set of depots. T is the number of days in the planning period, and Ω_{\max} is the maximum volume that should or can be delivered in the planning period. To obtain Ω_{\max} , the minimum of the average expected demand over the next n planning periods, and the total expected selected volume ($\Omega_{seed} + \Omega_{clustered}$) is calculated. We use the average expected demand over the next n planning periods, to balance the workload over a long period of time. We use the minimum of the two terms, since we can not plan more volume than we have available in our customer selection.

$$\Omega_{\max} = \max(\Omega_{mustgo}, \min(\Omega_{forecasted}, (\Omega_{seed} + \Omega_{clustered}))) \quad (G.2)$$

Additionally, the maximum of the minimum outcome and the expected delivery volume of all ‘must-go’ customers is Ω_{\max} , as illustrated in Equation G.2. If the forecasted volume is not sufficient to plan all ‘must-go’ customers, the ILP is infeasible, since the workload constraint can not be satisfied. Therefore, the expected volume for all deliveries to ‘must-go’ customers is added, and the maximum of these two is chosen for the workload balance.

To ensure feasibility, we have to calculate the expected volume for a certain day per depot with Equation G.3.

$$v_{\text{exp } tp} = \left[\left(\frac{1}{D_t^{\text{total}}} \right) \sum_{c \in C} \left(\left(\frac{\sum_{s \in S_c} m_{st}}{S_c^{\text{total}}} \right) \cdot r_{st} \right) \right] \quad (\text{G.3})$$

Where m_{st} is 1 if there is a delivery on day t in schedule s , and 0 if there is no delivery on day t in schedule s . S_c^{total} is the number of schedules that are allowed for customer c . Customer c is in the total customer set C , thus it includes seed customers as well as clustered customers. This is done, since we want to include the seed customers in balancing the volume equally, the seed volume is subtracted from the workload in later stage. D_t^{total} is the number of opened depots on day t . r_{st} is the volume that is planned in schedule s on day t , calculated by multiplying the relative delivery size with the customers capacity. The result of Equation G.3 is the volume that equals the total expected delivery volume for day t , and depot p .

The maximum delivery volume per depot, day, and vehicle type set is the maximum of Equations G.1 and G.3:

$$v_{\text{max } tp} = \max(v'_{\text{max } tp}, v_{\text{exp } tp}) \quad (\text{G.4})$$

Now that we have calculated the workload per depot and per day, we can easily transfer this to a workload per seed delivery, since we know the assignment of the seed delivery to depots. We divide the workload per depot and day by the number of seed deliveries assigned to this depot on this day. Additionally, we subtract the delivery volume of the specific seed delivery to get the workload per seed delivery.

After calculating all $v_{\text{max } tp}$ for all days and p , we can select the appropriate maximum workload for every seed delivery, by dividing $v_{\text{max } tp}$ by the number of seeds that is assigned to day t , and depot d . The customer seed delivery volume is subtracted, since that is already known. By doing so, we get $v_{\text{max } k^p}$, which is the maximum workload for seed delivery k^p , a seed delivery k assigned to day t and depot p .

G.2 Minimum workload

Next to a maximum workload, a minimum workload is added to the ILP to find a well-balanced plan. The calculation of the minimum workload is based on the same principles as the maximum workload. Equation G.5 illustrates this.

$$v_{\min tp} = (1 - \beta) \cdot \left(\frac{VC_{tp}}{\sum_{j \in D} \sum_{i=1}^T VC_{ij}} \right) \cdot \Omega_{\max} \quad (\text{G.5})$$

When the expected delivery volume, calculated with Equation G.3 is below the outcome of Equation G.5 for a day and a depot, the minimum workload for this day and depot is set to the value of Equation G.3. This is done to ensure the ILP is feasible, and the minimum workload constraint can be met. Equation G.6 illustrates this.

$$v_{\min tp} = \min(v'_{\min tp}, v_{\exp tp}) \quad (\text{G.6})$$

G.3 Improving the balance of the workloads

When the minimum (or maximum) workload for a certain day and depot is calculated with Equation G.2, and not with Equation G.4 (or G.1), the total minimum (or maximum) workload does not equal the summation of $w'_{\min tp}$ over all days t , and depots p anymore. We have to correct all the workload calculations for the other days and depots, such that the total workload is equal to this summation again. This problem is also apparent in the assignment of clustered customers to delivery schedules, and Appendix H explains the method that is used for this procedure.

Appendix H: Workload calculation: Improving the balance

After calculating the workload settings for every day, we need to make sure the complete volume is balanced. Therefore, the sum of all workloads per day and depot must equal the total workload calculated for the entire planning period. This is the total number of seed deliveries for the assignment of seed customers times $(1 + \beta_{seed})$, and Ω_{max} times $(1 + \beta)$ for the assignment of clustered customers.

We explain this with an example for improving the balance of the workloads for the assignment of clustered customers. The sum of all $w_{max tp}$ must equal $(1 + \beta) \cdot \Omega_{max}$, and the sum of all $w_{min tp}$ must equal $(1 - \beta) \cdot \Omega_{max}$. To do this, we evaluate every $w_{max tp}$ and $w_{min tp}$, and see if the expected volume of ‘must-go’ customers for a certain day and depot is not higher than $w_{max tp}$ or that the total volume of ‘must-go’, and ‘may-go’ customers is not lower than $w_{min tp}$. If that is the case, we reset $w_{max tp}$ or $w_{min tp}$ for a certain day t , and recalculate all other $w_{max tp}$ or $w_{min tp}$. By subtracting the amount that was just set for day t from the total workload amount, and recalculating with the new workload, we eventually find the best achievable workload balance. Table I.1 illustrates an example of this method.

Day	Depot	Calculated workload	Volume ‘must-go’ per depot	Percentage of vehicle capacity	New workload
1	1	1000	1200	12,50%	1200
1	2	1000	1200	12,50%	1200
2	1	500	300	6,25%	400
2	2	1500	300	18,75%	1200
3	1	1000	1100	12,50%	1100
3	2	1000	1100	12,50%	1100
4	1	1000	900	12,50%	900
4	2	1000	900	12,50%	900
Total	-	8000	7000	100%	8000

Table H.1. An example of the calculation procedure of the maximum workload for a day and a depot.

$$w_{max tp} = (1 + \beta) \cdot \left(\frac{VC_{tp}}{\sum_{j \in D} \sum_{i=1}^T VC_{ij}} \right) \cdot \Omega_{max} \quad (H.1)$$

The calculated workload is based on Equation I.1, which is defined in Appendix G. For the ILP to be feasible, the maximum workload of Days 1 and 3 need to be raised, the total raise is 600, this is subtracted from the workload on Days 2 and 4, by dividing the subtraction equally over all the available days and depots from which a subtraction still can take place, according to the relation of their vehicle capacity. If we calculate it for Day 4, we get 850 litres $((12,50 / (12,50 + 12,50 + 18,75 + 6,25)) * 600 = 150; 1000 - 150 = 850)$, thus the new workload is set to 900 litres for both depots on Day 4, since the ‘must-go’ volume is larger than this 850 litres. The total raise is still 400 litres. The new workload for Day 2 and Depot 1 can be calculated with $(6,25 / (18,75 + 6,25)) * 400 = 100; 500 - 100 = 400)$. With this procedure, feasibility of the ILPs with regards to workload is guaranteed.

Appendix I: Alternative designs: Balance customers

Balance customers are designed to cope with route failures, which occur in gas distribution. A route failure is a difference between the expected delivery volume in a trip, and the actual delivery volume. A route failure results in either too little or too many customers in the pre-planned trips. The problem decreases utilization and this leads to a lower volume per kilometre measure. Therefore, an alternative design that includes balance customers that are added as the last stop in a trip with a flexible delivery volume. This flexible delivery volume is given by a minimum and a maximum delivery. The minimum delivery is given by the minimum delivery for which a delivery is still economically feasible. The maximum delivery is given by the maximum delivery that can still fit in the balance customer's tank. A balance customer is always a 'may-go' customer, thus there is no problem if we can not deliver to the customer due to a route failure. Figure I.1 illustrates the alternative design of the algorithm with balance customers.

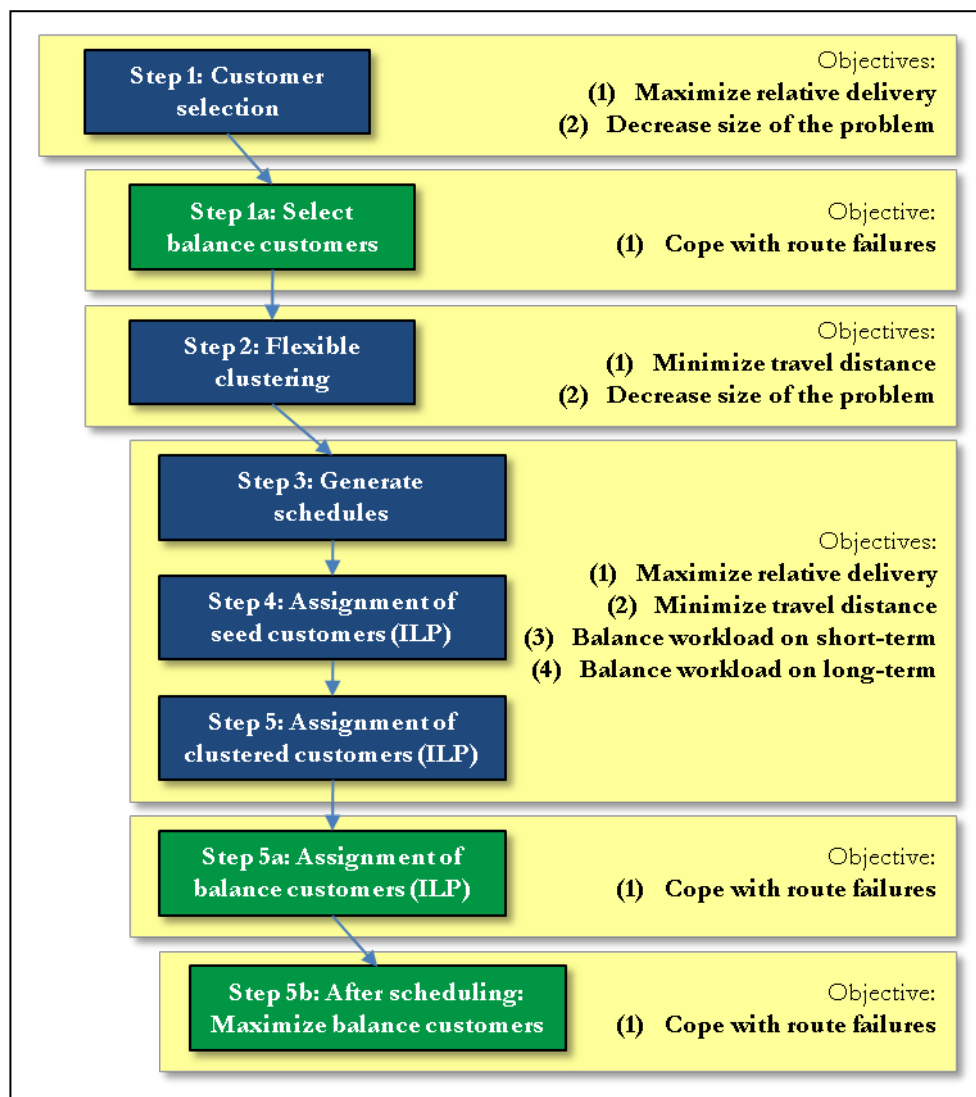


Figure I.1 An illustration of the alternative design of the algorithm with balance customers. The green blocks are the newly added steps.

Step 1a: Select balance customers

After selecting set of customers that is considered in the algorithm, we select the balance customers in order to cope with the problem of route failures (Dror and Ball, 1987; Dror and Trudeau, 1988). To ensure that the last stop in a trip is not of high cost, the balance customers are selected around the depots, this means the customer can be easily added as the last customer in a trip, since the vehicle has to return to the depot at the end of a trip.

Step 5a: Schedule assignment with an ILP for balance customers

The assignment of a balance customer to a schedule and a seed delivery is very similar to the assignment of a clustered customer, but the workload is calculated differently.

Workload constraint

A balance customer is always the last customer in a scheduled trip if it is in the trip. We assign a single balance customer to a seed customer delivery, since we do not know how many trips we will have after scheduling. Therefore, we use the number of seed deliveries as an indication for the number of trips. We could also use the number of vehicles available, but as we said before, we do not have exact information about the number of trips these vehicles will perform.

Vehicle capacity

If there is no vehicle available at any of the depots that can serve customer c on the day of delivery t in schedule s that is connected to customer c , the combination of schedule s and customer c is excluded from the analysis. If there is no possible schedule left for customer c , the schedule will be allowed to ensure the ILP is feasible. The customer is then assigned to a day where there is no vehicle capacity for the customer.

Figure I.2 illustrates the ILP that is used to assign the balance customers to a schedule and a seed customer delivery. We use the assigned seed customers in Step 4 as a reference point for the distance calculations in assigning the balance customers to schedules.

$$\min \sum_{k^p \in SCD} \sum_{c \in BC} \left(\left(\sum_{s \in S_c} d_{ck^p} \cdot X_{sck^p} \right) + \left(\frac{\tilde{y}_c^3}{10^4 \cdot \sqrt[4]{\tilde{d}_c}} \right) X_c \right) \quad (I.1)$$

s.t.

$$\sum_{k^p \in SCD} \sum_{s \in S_c} X_{sck^p} + X_c = 1 \quad \text{where } c \in BC \quad (I.2)$$

$$\sum_{c \in BC} \sum_{s \in S_c} X_{sck^p} \leq 1 \quad \text{where } t = 1, \dots, T; k^p \in SCD \quad (I.3)$$

$$X_{sck^p} \leq r_{st} \quad \text{where } t = 1, \dots, T, k^p \in SCD, c \in BC \quad (I.4)$$

$$X_{sck^p} \in \{0, 1\} \quad \text{where } s \in S_c, c \in BC, k^p \in SCD \quad (I.5)$$

$$X_c \in \{0, 1\} \quad \text{where } c \in BC \quad (I.6)$$

Variables

X_{sck^p} binary variable that equals 1 if delivery schedule s is assigned to customer c , and seed delivery k^p , 0 otherwise

X_c binary variable that equals 1 if customer c is excluded from the current planning horizon, 0 otherwise

Data

\tilde{y}_c relative delivery size in percentage if customer c is forwarded

d_{ck^p} distance contribution to assign customer c to seed delivery k^p

\tilde{d}_c is the distance from customer c to its nearest allowed seed delivery

r_{st} integer variable for the delivery volume on day t in schedule s

k^p seed delivery k that is assigned to day t and depot p

S_c set of allowable delivery schedules for customer c

SCD set of seed customer deliveries

BC set of balance customers

T number of days in the planning horizon

Figure I.2. The ILP to assign the balance customers to schedules, while satisfying the constraints and balancing the number of balance customers over the planning period.

The ILP has the objective to minimize the additional distance of adding a customer to a seed. The cost for forwarding a balance customer is the same as for a clustered customer. Constraint I.2 is added to ensure a schedule is chosen, or the customer is forwarded to a future planning period. Constraint I.3 ensures that every seed customer delivery has only one balance customer assigned. Constraint I.4 ensures a schedule can not be assigned a balance customer and a seed customer delivery if there is no delivery on the same day as the seed customer delivery in schedule s for customer c . Constraints I.5 and I.6 ensure the decision variables satisfy the binary properties. For the ILP for balance customers, we can use the same distance types as in the ILP for clustered customers.

Step 5b: Maximize balance customers

The balance customers are scheduled at the end of a trip. The actual delivery volume will be somewhere between the volume that can be delivered to the customer, and the volume that is generally applied as a minimum for a delivery. The use of balance customers in this way ensures a high utilization of vehicles, and ensures a full vehicle can leave the depot for a trip. Additionally, the balance customer serves as a buffer in actually driving the trip. This buffer can be used when route failures occur (Dror and Trudeau, 1988). If there are too many customers in the trip, we just skip the balance customer, since it is a ‘may-go’ customer. If there are too little customers in the trip, we can deliver up to the balance customer’s maximum delivery. Figure I.3 illustrates this idea, where the green colour signifies the balance customer that can receive a volume between 200 litres and 850 litres. Eventually, the customer is calculated to receive 300 litres, since that fills up the vehicle entirely. The algorithm can have more balance customers per trip, but we use one balance customer per trip, because that already illustrates the idea of using the last customer as a buffer for the complete trip.

Variable	Volume
Minimum economic delivery size	200
Vehicle size	2000

Table I.1. The variables for the example of maximizing the balance customers.

Customer	Volume
Customer 1	500
Customer 2	300
Customer 3	400
Customer 4	500
Balance customer	200 - 850

Table I.2. The customers in the example for maximizing the balance customers.

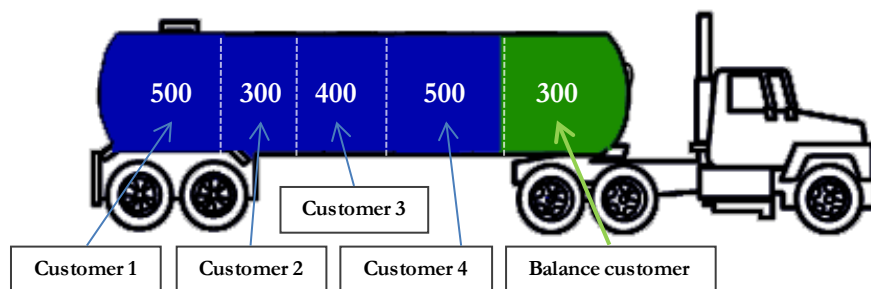


Figure I.3. The figure illustrates how the last customer in the vehicle is the balance customer with the flexible delivery volume. The delivery volume to the balance customer in this example is 300.

Appendix J: Alternative designs: Different objective functions

We test six different set-ups for the objective functions in the ILP for assigning clustered customers. We test an objective function that simply divides the relative delivery size by the distance measure. This is a combination of the volume per kilometre measure, and the relative delivery. The objective function maximizes the relative delivery per kilometre. Equation J.1 illustrates this, and the negative sign for forwarding a ‘may-go’ customer ensures that we plan as many ‘may-go’ customers as possible.

$$\max \sum_{k^{tp} \in SCD} \sum_{c \in BC} \left(\left(\sum_{s \in S_c} \sum_{t=0}^T \left(\frac{y_{ts}}{d_{ck^{tp}}} \right) X_{sck^{tp}} \right) - \left(\frac{\tilde{y}_c}{\tilde{d}_c} \right) X_c \right) \quad (J.1)$$

Where y_{ts} is the relative delivery size of schedule s on day t , and $d_{ck^{tp}}$ is the distance between the customer c and the seed delivery k^{tp} . \tilde{y}_c is the relative delivery if the customer c is forwarded to a future planning period, and \tilde{d}_c is the distance to the nearest seed delivery. $X_{sck^{tp}}$ is the decision variable that assigns a customer c to a schedule s , and a seed delivery k^{tp} . X_c is the decision variable that forwards a customer c .

Additionally we test an objective function with a bi-criteria approach, in which we can influence the importance of either relative delivery or distance by setting parameters λ and μ respectively. Equation J.2 illustrates the bi-criteria objective function, and the negative sign for forwarding a ‘may-go’ customer ensures that we plan as many ‘may-go’ customers as possible.

$$\max \sum_{k^{tp} \in SCD} \sum_{c \in BC} \left(\left(\sum_{s \in S_c} \sum_{t=0}^T \left(\lambda \cdot y_{ts} + \mu \cdot \frac{1}{d_{ck^{tp}}} \right) X_{sck^{tp}} \right) - \left(\lambda \cdot \tilde{y}_c + \mu \cdot \frac{1}{\tilde{d}_c} \right) X_c \right) \quad (J.2)$$

Additionally, these objective functions are tested with a positive value for forwarding a ‘may-go’ customer to the next period, because in the Equations J.1 and J.2, the ILP will always plan the maximum volume. Equations J.3 and J.4 are the objective functions in Equations J.1 and J.2 with a positive forwarding value.

$$\max \sum_{k^{tp} \in SCD} \sum_{c \in BC} \left(\left(\sum_{s \in S_c} \sum_{t=0}^T \left(\frac{y_{ts}}{d_{ck^{tp}}} \right) X_{sck^{tp}} \right) + \left(\frac{\tilde{d}_c}{\tilde{y}_c} \right) X_c \right) \quad (J.3)$$

$$\max \sum_{k^{tp} \in SCD} \sum_{c \in BC} \left(\left(\sum_{s \in S_c} \sum_{t=0}^T \left(\lambda \cdot y_{ts} + \mu \cdot \frac{1}{d_{ck^{tp}}} \right) X_{sck^{tp}} \right) + \left(\kappa \cdot (100 - \tilde{y}_c) + \tilde{d}_c \right) X_c \right) \quad (J.4)$$

The forwarding expression in Equation J.3 differs from the expression in Equation J.2 in that the distance and relative delivery size are exchanged. The forwarding expression in Equation J.4 differs from the expression in Equation J.2 in that κ now is the parameter to set the importance of the relative delivery size. Furthermore, the relative delivery size is subtracted from 100, and the distance measure is not divided by μ anymore. These changes are necessary, because we want to push ‘may-go’ customers that have a high

distance to their nearest delivery, or a relatively low relative delivery size, to a following planning period.

Additionally, we have created an objective function based on Equation J.4, but with a more complex forwarding function. This complex forwarding expression has the objective to have an exponential growth of the value for forwarding a ‘may-go’ customer when the relative delivery gets lower, or the distance to the nearest seed customer delivery gets higher. Equation J.5 illustrates this objective function, and it is found by simulating the relative delivery size and the distance to the nearest seed customer delivery for a small set of customers.

$$\max \sum_{k^{tp} \in SCD} \sum_{c \in BC} \left(\left(\sum_{s \in S_c} \sum_{t=0}^T \left(\lambda \cdot y_{ts} + \mu \cdot \frac{1}{d_{ck^{tp}}} \right) X_{sck^{tp}} \right) + \left(\frac{(100 - \tilde{y}_c)^4}{10^4} + \frac{\tilde{d}_c^2}{10^2} \right) X_c \right) \quad (J.5)$$

We also test the objective function used in the Period Scheduler, and in Step 5 of the algorithm in this thesis. It is given in Equation J.6.

$$\min \sum_{k^{tp} \in SCD} \sum_{c \in BC} \left(\left(\sum_{s \in S_c} d_{ck^{tp}} \cdot X_{sck^{tp}} \right) + \left(\frac{\tilde{y}_c^3}{10^4 \cdot \sqrt[4]{\tilde{d}_c}} \right) X_c \right) \quad (J.6)$$

Appendix K: Alternative designs: ILP with combination check

$$\min \sum_{k^p \in SCD} \sum_{c \in CC} \left(\left(\sum_{s \in S_c} d_{ck^p} \cdot X_{sck^p} \right) + \left(\frac{\tilde{y}_c^3}{10^4 \cdot \sqrt[4]{\tilde{d}_c}} \right) X_c \right) \quad (K.1)$$

s.t.

$$\sum_{k^p \in SCD} \sum_{s \in S_c} X_{sck^p} + X_c = 1 \quad \text{where } c \in CC \quad (K.2)$$

$$v_{\min k^p} \leq \sum_{c \in CC} \sum_{s \in S_c} r_{st} X_{sck^p} + \sum_{c \in CC} \sum_{s \in S_c} \sum_{g \in SCD} R_{sk^p} X_{scg} \leq v_{\max k^p} \quad (K.3)$$

where $t = 1, \dots, T$; $k^p \in SCD$

$$X_{sck^p} \leq r_{st} \quad \text{where } t = 1, \dots, T, k^p \in SCD, c \in CC \quad (K.4)$$

$$X_{sck^p} \leq \varsigma_{k^p c} \quad \text{where } k^p \in SCD, c \in CC \quad (K.5)$$

$$X_c = 0 \quad \text{where } c \in CC_{mustgo} \quad (K.6)$$

$$X_{sck^p} \in \{0, 1\} \quad \text{where } s \in S_c, c \in CC, k^p \in SCD \quad (K.7)$$

$$X_c \in \{0, 1\} \quad \text{where } c \in CC \quad (K.8)$$

Variables

X_{sck^p} binary variable that equals 1 if delivery schedule s is assigned to customer c , and seed delivery k^p , 0 otherwise

X_c binary variable that equals 1 if customer c is excluded from the current planning horizon, 0 otherwise

Data

y_{ts} relative delivery size in percentage on day t in schedule s

\tilde{y}_c relative delivery size in percentage if customer c is forwarded

d_{ck^p} distance contribution to assign customer c to seed delivery k^p

\tilde{d}_c is the distance from customer c to its nearest allowed seed delivery

k^p seed delivery k that is assigned to day t , and depot p

$v_{\max k^p}$ maximum workload for seed delivery k^p on day t

$v_{\min k^p}$ minimum workload for seed delivery k^p on day t

r_{st} the delivery volume as an integer on day t in schedule s

$\varsigma_{k^p c}$ is 1 if the customer c can be assigned to seed delivery k^p , 0 otherwise

S_c set of allowable delivery schedules for customer c

R_{sk^p} volume after the first delivery in schedule s that is assigned to seed k^p

SCD set of seed customer deliveries

CC set of clustered customers

CC_{mustgo} set of clustered customers that are ‘must-go’ customers

T number of days in the planning horizon

Figure K.1. The ILP to assign the clustered customers to schedules, while satisfying the customer combination constraint in Equation K.5. We set $\varsigma_{k^p c}$ to 1 if a clustered customer c can be assigned to seed customer delivery k^p , and we set it to 0 otherwise.

Appendix L: Settings used in the computational tests

Setting	Value
α (the minimum relative delivery)	0,40
Number of weeks forecasted volume	13 weeks
Distance type clusters	Real distance
Minimum cluster size	0
Maximum cluster size	15
Threshold distance cluster	15 km
Minimum number of seeds per day	5
β (bandwidth for clustered customers)	0,05
Solution procedure ILP seed customers	Branch-and-bound
Solution procedure ILP clustered customers	Rounding
Distance type ILP seed customers	Real distance
Distance type ILP clustered customers	Detour distance
Schedule algorithm	Insertion

Table L.1. Settings used to compare the actually driven plans with the calculated plans of the algorithm.

Appendix M: Dataset overview

Vehicle	Capacity (litres)	Start time (hrs)	End time (hrs)	Lunch break (mins)	Depot
CL04	11.764	06:00	18:00	45	Depot 2
CL05	16.823	06:00	18:00	45	Depot 2
CL06	12.352	06:00	18:00	45	Depot 2
NW02	17.058	06:00	18:00	45	Depot 3
NW03	11.882	06:00	18:00	45	Depot 3
NW05	12.353	06:00	18:00	45	Depot 3
PE06	14.117	06:00	18:00	45	Depot 2
PE08	9.411	06:00	18:00	45	Depot 1
SW05	12.352	06:00	18:00	45	Depot 1
SW06	16.470	06:00	18:00	45	Depot 1

Figure M.1. Vehicles available in the Northern-England region (NE-region).

Depot	Gates	Opening times (hrs)
Depot 1	1	06:00 - 22:00
Depot 2	1	06:00 - 22:00
Depot 3	1	06:00 - 22:00

Figure M.2. Depots available in the NE-region.

Description	Fictive cost	Working time
Unloading at a customer (fixed cost)	-	5 minutes
Unloading at a customer (variable cost)	-	1 minute per 1.000 litres
Loading at the depot (fixed cost)	-	20 minutes
Loading at the depot (variable cost)	-	-
Per driven kilometre (variable cost)	1	-
Per worked minute (variable cost)	1	-
Per worked minute overtime (variable cost)	2	-
Per vehicle per day (fixed cost)	100	-
Per trip (fixed cost)	500	-

Figure M.3. Fictive costs in the NE-region.

Appendix N: Comparison actual routes and SHORTREC routes

This appendix illustrates the difference between the actually driven routes and the routes SHORTREC would have calculated. The actually driven routes are scheduled by the planners, and the kilometres for the actually driven routes are calculated with SHORTREC. Additionally, we let SHORTREC calculate routes with an insertion heuristic described in Poot et al. (2002). We use the vehicles that are used by the planners as the available vehicles for the SHORTREC planning. Table N.1 illustrates the significant gains that SHORTREC can provide to the planners. In the SHORTREC plans, constraints such as vehicle capacity, time windows, and allowed vehicle types are checked, which are not checked by the planners.

Scheduling Method	Num. of orders	Volume (litres)	Dist. (km)	Volume/kilometre (litres/km)	Cost (fictive)	Cost/Litre	Veh. used	Veh. Utiliz.	Work Time (hrs)	Volume/kilometre Improve.	Cost/Volume Savings
Year 2007 Week 49											
Actual	378	388.162	7.237	53,6358	23.181	0,0597	35	76,57%	207,40	-	-
SHORTREC	378	388.162	6.279	61,8191	20.358	0,0524	29	79,05%	186,32	15,26%	12,18%
Year 2007 Week 50											
Actual	375	383.121	7.659	50,0223	23.874	0,0623	35	71,41%	208,51	-	-
SHORTREC	364	373.272	6.305	59,2025	20.124	0,0539	28	78,27%	183,65	18,35%	13,48%

Table N.1. The differences between the resulting volume per kilometre, vehicles used, and costs per litre indicate that the use of SHORTREC provides great wins.




Figure N.1. The results in a graph illustrate the benefits of using SHORTREC.

Appendix O: Computational results: The weekly comparison

Setting	Num. of orders	Volume (litres)	Dist. (km)	Volume/ distance (litres/km)	Cost (fictive)	Cost/ litre	Veh. used	Veh. utilization	Average relative delivery	Volume/ distance improve.	Cost/ litre savings
Week 40											
Base scenario	212	226.248	4.236	53,4108	13.286	0,0587	19	64,15%	53,56%	-	-
Algorithm	311	319.827	4.860	65,8080	16.027	0,0501	23	69,62%	60,30%	23,21%	14,66%
Week 41											
Base scenario	217	257.604	4.784	53,8470	14.622	0,0568	19	74,65%	53,29%	-	-
Algorithm	316	327.388	5.022	65,1908	16.444	0,0502	22	68,14%	59,29%	21,07%	11,51%
Week 42											
Base scenario	219	234.605	4.428	52,9822	13.853	0,0590	18	71,11%	53,80%	-	-
Algorithm	351	334.727	5.213	64,2101	17.273	0,0516	23	71,16%	59,74%	21,19%	12,61%

Table O.1. The individual results separated by week.




Figure O.1. The volume per kilometre measures for the algorithm, for weeks 40, 41, and 42, are better than the base scenario.




Figure O.2. The average relative delivery results for the algorithm in weeks 40, 41, and 42, are better than the base scenario.

Appendix P: Computational results: The working time balance

The working time balance for the algorithm plans are not perfect, but especially week 42 shows a connection to the workload balance.




Figure P.1. The working time balance for the algorithm plans.




Figure P.2. The working time balance for the actual plans.

Appendix Q: Computational results: A different vehicle set-up

This appendix illustrates the separate weekly results and the geographical results for week 42, for the algorithm where the vehicle capacities are set-up differently. Saturday has 50% of the vehicle capacity of the other days in the planning period. The geographical results illustrate that there is workload at every depot on every day of the planning period.

Scenario	Num. of orders	Volume (litres)	Dist. (km)	Volume/distance (litres/km)	Cost (fictive)	Cost/litre	Veh. used	Veh. utilization	Average relative delivery	Volume/distance improve.	Cost/volume savings
Week 40											
Base scenario	212	226.248	4.236	53,4108	13.286	0,0587	19	64,15%	53,56%	-	-
Algorithm: Saturday 100%	310	319.827	4.860	65,8080	16.027	0,0501	23	69,62%	60,30%	23,21%	14,66%
Algorithm: Saturday 50%	312	312.986	4.805	65,1376	15.883	0,0507	23	73,57%	60,04%	21,96%	13,58%
Week 41											
Base scenario	217	257.604	4.784	53,8470	14.622	0,0568	19	74,65%	53,29%	-	-
Algorithm: Saturday 100%	316	327.388	5.022	65,1908	16.444	0,0502	22	68,14%	59,29%	21,07%	11,51%
Algorithm: Saturday 50%	316	324.166	5.021	64,5620	16.347	0,0504	21	72,73%	59,14%	19,90%	11,16%
Week 42											
Base scenario	219	234.605	4.428	52,9822	13.853	0,0590	18	71,11%	53,80%	-	-
Algorithm: Saturday 100%	351	334.727	5.213	64,2101	17.273	0,0516	23	71,16%	59,74%	21,19%	12,61%
Algorithm: Saturday 50%	354	324.638	4.925	65,9163	16.451	0,0507	22	72,57%	59,33%	24,41%	14,18%

Table Q.1. The separate weekly results of setting the vehicle capacity on Saturday to 50%.



Figure Q.2. The geographical results for setting the vehicle capacity on Saturday to 50% in week 42. Note that one customer block can exist of several customers.

Appendix R: Computational results: α parameter

α	Volume	Distance	Average relative delivery	Algorithm run time (seconds)	Volume/kilometre (litres/km)	Number of selected customers
0,35	327.415	5.464	57,56%	1.854	59,9222	630
0,40	327.388	5.022	59,29%	1.183	65,1908	499
0,45	307.267	4.860	60,84%	492	63,2237	401
0,50	253.169	4.701	63,48%	107	53,8543	301
0,55	232.923	4.908	64,62%	29	47,4578	223

Figure R.1. The α setting results for week 41 indicate the increasing performance measure of average relative delivery, but show the weaker results for the volume per kilometre measure.

α	Volume (litres)	Distance (km)	Average relative delivery	Algorithm run time (seconds)	Volume/kilometre (litres/km)	Number of selected customers
0,35	318.772	5.137	58,65%	2.996	62,0541	678
0,40	334.727	5.213	59,74%	1.181	64,2101	530
0,45	304.968	5.272	61,47%	542	57,8467	417
0,50	259.014	4.957	63,79%	166	52,2522	312
0,55	212.710	4.282	65,29%	58	49,6754	237

Figure R.2. The α setting results for week 42 clearly indicate the strongly decreasing run time, but also decreasing volume per kilometre performance measure.

Appendix S: Computational results: Bandwidth parameter

Bandwidth	Num. of orders	Volume (litres)	Dist. (km)	Volume/ kilometre (litres/km)	Cost (fictive)	Cost/ litre	Veh. used	Veh. utilization	Average relative delivery	Volume/ kilometre improve.	Cost/ volume savings
Base scenario*	219	234.605	4.428	52,9822	13.853	0,0590	18	71,11%	53,80%	-	-
0,01	347	328.150	5.083	<u>64,5583</u>	16.949	0,0517	23	70,92%	<u>59,77%</u>	21,85%	12,53%
0,02	347	328.974	5.078	<u>64,7842</u>	16.873	0,0513	23	71,32%	<u>59,72%</u>	22,28%	13,14%
0,03	349	331.927	5.170	<u>64,2025</u>	17.208	0,0518	24	71,28%	<u>59,72%</u>	21,18%	12,20%
0,04	350	332.775	5.158	<u>64,5163</u>	17.154	0,0515	23	70,70%	<u>59,68%</u>	21,77%	12,70%
0,05	351	334.727	5.213	<u>64,2101</u>	17.273	0,0516	23	71,16%	<u>59,74%</u>	21,19%	12,61%
0,06	351	334.826	5.201	<u>64,3772</u>	17.329	0,0518	24	68,99%	<u>59,76%</u>	21,51%	12,35%
0,07	352	335.430	5.193	<u>64,5927</u>	17.294	0,0516	24	70,42%	<u>59,67%</u>	21,91%	12,69%
0,08	352	335.359	5.234	<u>64,0732</u>	17.339	0,0517	23	69,40%	<u>59,67%</u>	20,93%	12,44%
0,09	355	326.992	5.054	<u>64,6996</u>	16.938	0,0518	24	71,19%	<u>59,61%</u>	22,12%	12,28%
0,10	354	323.616	5.048	<u>64,1078</u>	16.883	0,0522	23	71,43%	<u>59,68%</u>	21,00%	11,65%

* The base scenario was planned by the planners, and therefore has no such setting.

Table S.1. The results show that the volume per kilometre, and the average relative delivery are almost constant when the bandwidth changes.




Figure S.1. The graph illustrates the consistent results of volume per kilometre and average relative delivery.

Appendix T: Computational results: Distance type parameter

All the separate results for detour distance are better on volume per kilometre in every week. The average relative delivery is better if we use real distance (km). This appendix shows that detour distance performs stable in all three weeks. The geographical results for real distance and detour distance are given in Figures S.1 and S.2, and they illustrate that the workload is spread out over the days and the depots.

Distance type	Num. of orders	Volume (litres)	Dist. (km)	Volume/ kilometre (litres/km)	Cost (fictive)	Cost/ litre	Veh. used	Veh. utilization	Average relative delivery	Volume/ kilometre improve.	Cost/ volume savings
Base scenario*	212	226.248	4.236	53,4108	13.286	0,0587	19	64,15%	53,56%	-	-
Real distance (km)	274	312.649	5.170	60,4737	16.554	0,0529	23	69,74%	61,70%	13,22%	9,84%
Real distance (time)	345	328.080	6.627	49,5066	20.760	0,0633	25	69,99%	59,76%	-7,31%	-7,76%
Euclidean distance	276	314.038	5.251	59,8054	16.806	0,0535	24	74,79%	61,48%	11,97%	8,87%
Detour distance	311	319.827	4.860	65,8080	16.027	0,0501	23	69,62%	60,30%	23,21%	14,66%

* The base scenario was planned by the planners, and therefore has no such setting.

Table T.1. The individual results for week 40 for different distance types.

Distance type	Num. of orders	Volume (litres)	Dist. (km)	Volume/ kilometre (litres/km)	Cost (fictive)	Cost/ litre	Veh. used	Veh. utilization	Average relative delivery	Volume/ kilometre improve.	Cost/ volume savings
Base scenario*	217	257.604	4.784	53,8470	14.622	0,0568	19	74,65%	53,29%	-	-
Real distance (km)	293	316.114	5.170	61,1439	16.472	0,0521	23	73,26%	60,06%	13,55%	8,20%
Real distance (time)	361	344.222	6.607	52,0996	20.708	0,0602	25	77,90%	59,08%	-3,25%	-5,99%
Euclidean distance	312	327.358	5.357	61,1085	17.268	0,0527	24	70,72%	59,17%	13,49%	7,07%
Detour distance	316	327.388	5.022	65,1908	16.444	0,0502	22	68,14%	59,29%	21,07%	11,51%

* The base scenario was planned by the planners, and therefore has no such setting.

Table T.2. The individual results for week 41 for different distance types.

Distance type	Num. of orders	Volume (litres)	Dist. (km)	Volume/ kilometre (litres/km)	Cost (fictive)	Cost/ litre	Veh. used	Veh. utilization	Average relative delivery	Volume/ kilometre improve.	Cost/ volume savings
Base scenario*	219	234.605	4.428	52,9822	13.853	0,0590	18	71,11%	53,80%	-	-
Real distance (km)	320	322.352	5.188	62,1342	16.931	0,0525	22	70,37%	60,63%	17,27%	11,05%
Real distance (time)	400	346.487	6.306	54,9456	20.534	0,0593	28	77,49%	59,00%	3,71%	-0,36%
Euclidean distance	344	333.716	5.396	61,8451	17.702	0,0530	24	71,62%	60,00%	16,73%	10,17%
Detour distance	351	334.727	5.213	64,2101	17.273	0,0516	23	71,16%	59,74%	21,19%	12,61%

* The base scenario was planned by the planners, and therefore has no such setting.

Table T.3. The individual results for week 42 for different distance types.



Figure T.1. The geographical results for week 42 when the real distance (km) is used.



Figure T.2. The geographical results for week 42 when the detour distance is used.

Appendix U: Computational results: Alternative designs

U.1 Balance customers

The balance customers are tested to illustrate how planners cope with route failures currently. We test the balance customers with a minimum economic delivery of 200 litres. Table U.1 illustrates that this method is inefficient, since the average relative delivery and the volume per kilometre are lower.

Balance customers	Volume (litres)	Dist. (km)	Volume/kilometre (litres/km)	Cost (fictive)	Cost/litre	Veh. used	Veh. utilization	Average relative delivery	Volume/kilometre improve.	Cost/volume savings
Week 40										
Base scenario*	226.248	4.236	53,4108	13.286	0,0587	19	64,15%	53,56%	-	-
Off	312.649	5.170	60,4737	16.554	0,0529	23	69,74%	61,70%	13,22%	9,84%
On	301.350	5.074	59,3910	16.252	0,0539	22	73,20%	61,55%	11,20%	8,16%
Week 41										
Base scenario*	257.604	4.784	53,8470	14.622	0,0568	19	74,65%	53,29%	-	-
Off	316.114	5.170	61,1439	16.472	0,0521	23	73,26%	60,06%	13,55%	8,20%
On	320.365	5.513	58,1108	17.444	0,0545	24	73,00%	59,51%	7,92%	4,07%
Week 42										
Base scenario*	234.605	4.428	52,9822	13.853	0,0590	18	71,11%	53,80%	-	-
Off	322.352	5.188	62,1342	16.931	0,0525	22	70,37%	60,63%	17,27%	11,05%
On	320.678	5.339	60,0633	17.069	0,0532	22	74,34%	60,98%	13,37%	9,86%

* The base scenario was planned by the planners, and therefore this setting is not applicable.

Table U.1. The table displays the results of using balance customers in the algorithm.

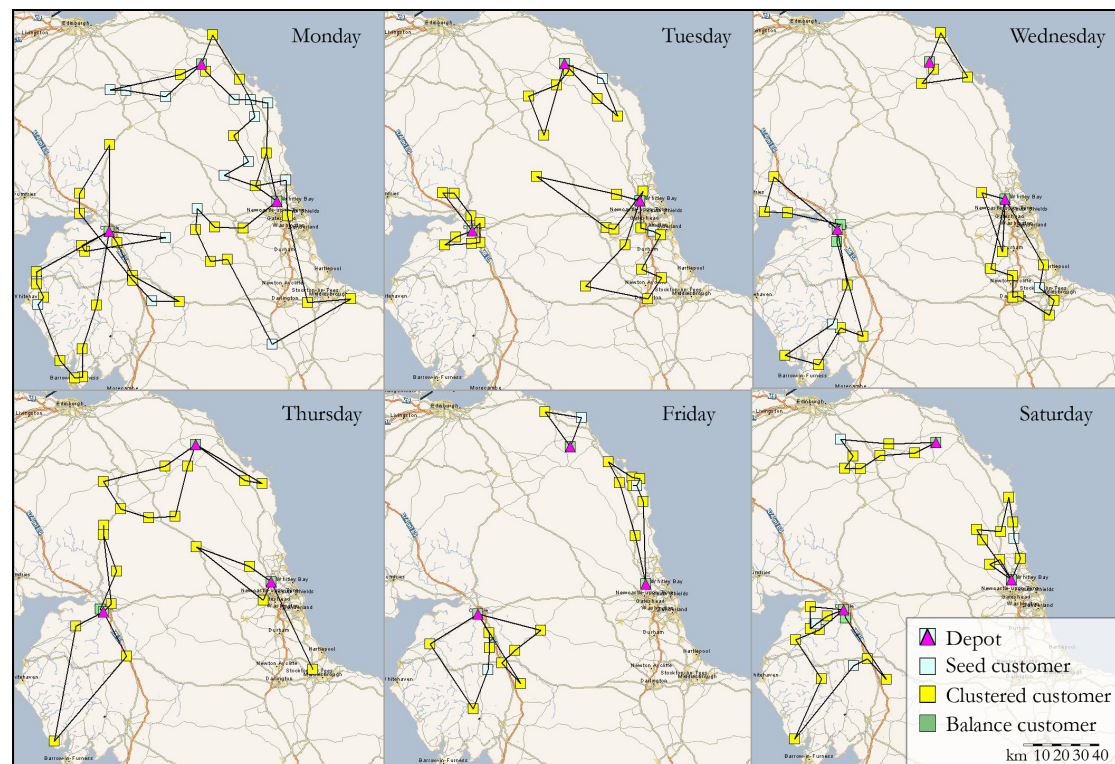


Figure U.1. The geographical spread in week 42 illustrates that the balance customers are in close proximity of the depots.

U.2 Objective functions

We test the different objective functions mentioned in Appendix J as Equations J.1 to J.6. Table U.2 shows the combined results of testing the different equations in the weeks 40, 41, and 42. We used real distance for these tests. Equation J.6 is the best objective function with **7,33%** savings in fictive costs per litre, and a **10,76%** increase in volume per kilometre. Even more interesting is the high average delivery of **60,65%**, against a 53.55% average delivery of the base scenario. This results in lower cost from a long-term perspective, since a lower number of total visits have to be made (Dror and Ball, 1987). The algorithm run time does not depend on the shape of the objective function.

Objective function	Setting	Num. of orders	Volume (litres)	Dist. (km)	Volume/ kilometre (litres/km)	Cost (fictive)	Cost/ litre	Veh. used	Average relative delivery	Volume/ kilometre improve.	Cost/ volume savings
Base scenario*	-	648	718.457	13.448	53,4248	41.761	0,0581	56	53,55%	-	-
J.1	-	1137	1.050.362	18.060	58,1596	58.257	0,0555	75	57,93%	8,86%	4,58%
J.2	$\lambda = 0.7$ $\mu = 400$	1174	1.050.100	18.459	56,8882	60.276	0,0574	80	57,56%	6,48%	1,25%
J.3	-	1040	1.042.757	18.167	57,3984	58.953	0,0565	78	57,17%	7,44%	2,74%
J.4	$\lambda = 0.7$ $\mu = 400$ $\nu = 4$	836	927.472	16.975	54,6375	52.969	0,0571	72	60,98%	2,27%	1,75%
J.5	$\lambda = 0.7$ $\mu = 400$	933	964.157	18.090	53,2978	57.158	0,0593	78	60,74%	-0,24%	-1,99%
J.6	-	884	942.393	15.926	59,1732	50.765	0,0539	68	60,65%	10,76%	7,33%

* The base scenario was planned by the planners, and therefore has no objective function.

Table U.2. The experimental results of the different objective functions tested for the weeks 40, 41, and 42.

U.3 Check for customer combinations

Table U.3 displays the results of testing this additional constraint. The figures show that when the additional constraint is left out, the algorithm obtains slightly better results.

Combination check	Num. of orders	Volume (litres)	Dist. (km)	Volume/ kilometre (litres/km)	Cost (fictive)	Cost/ litre	Veh. used	Average relative delivery	Volume/ kilometre improve.	Cost/ volume savings
Week 40										
Base scenario*	212	226.248	4.236	53,4108	13.286	0,0587	19	53,56%	-	-
Off	272	301.350	5.074	59,3910	16.252	0,0539	22	61,55%	11,20%	8,16%
On	277	302.415	5.082	59,5071	16.313	0,0539	24	60,97%	11,41%	8,14%
Week 41										
Base scenario*	217	257.604	4.784	53,8470	14.622	0,0568	19	53,29%	-	-
Off	303	320.365	5.513	58,1108	17.444	0,0545	24	59,51%	7,92%	4,07%
On	302	319.985	5.500	58,1791	17.334	0,0542	23	59,69%	8,05%	4,56%
Week 42										
Base scenario*	219	234.605	4.428	52,9822	13.853	0,0590	18	53,80%	-	-
Off	309	320.678	5.339	60,0633	17.069	0,0532	22	60,98%	13,37%	9,86%
On	312	321.349	5.384	59,6859	17.321	0,0539	22	60,76%	12,65%	8,72%

* The base scenario was planned by the planners, and therefore has no combination check.

Table U.3. The experimental results of leaving out the constraint for weeks 40, 41, and 42 are slightly better.

U.4 Vehicle type knowledge

Chapter 4 explained the use of knowledge about vehicle type restrictions to determine the total workload. Table U.4 displays the results of the tests with the use of knowledge about vehicle type restrictions, and without. The figures show that the use of this additional information obtains very unstable and worse results over time.

Vehicle restrictions	Num. of orders	Volume (litres)	Dist. (km)	Volume/kilometre (litres/km)	Cost (fictive)	Cost/ litre	Veh. used	Average relative delivery	Volume/kilometre improve.	Cost/volume savings
Week 40										
Base scenario*	212	226.248	4.236	53,4108	13.286	0,0587	19	53,56%	-	-
Off	272	301.350	5.074	59,3910	16.252	0,0539	22	61,55%	11,20%	8,16%
On	278	306.922	5.107	60,0983	16.314	0,0532	22	60,56%	12,52%	9,48%
Week 41										
Base scenario*	217	257.604	4.784	53,8470	14.622	0,0568	19	53,29%	-	-
Off	303	320.365	5.513	58,1108	17.444	0,0545	24	59,51%	7,92%	4,07%
On	278	296.978	5.341	55,6034	16.579	0,0558	22	59,94%	3,26%	1,65%
Week 42										
Base scenario*	219	234.605	4.428	52,9822	13.853	0,0590	18	53,80%	-	-
Off	309	320.678	5.339	60,0633	17.069	0,0532	22	60,98%	13,37%	9,86%
On	304	314.916	5.397	58,3502	17.197	0,0546	23	60,90%	10,13%	7,52%

* The base scenario was planned by the planners, and therefore this setting is not applicable.

Table U.4. The results of testing the use of vehicle type knowledge show that it is better to not use this knowledge, since the results are not stable and overall worse.

U.5 Period Scheduler seed selection method

We test the seed selection method used in the Period Scheduler. The balancing criterion is the number of customers in a region. We test one and two centres per depot per day in the planning period. Figure U.2 illustrates the seed selection procedure of the Period Scheduler versus the seeds points in flexible clustering.

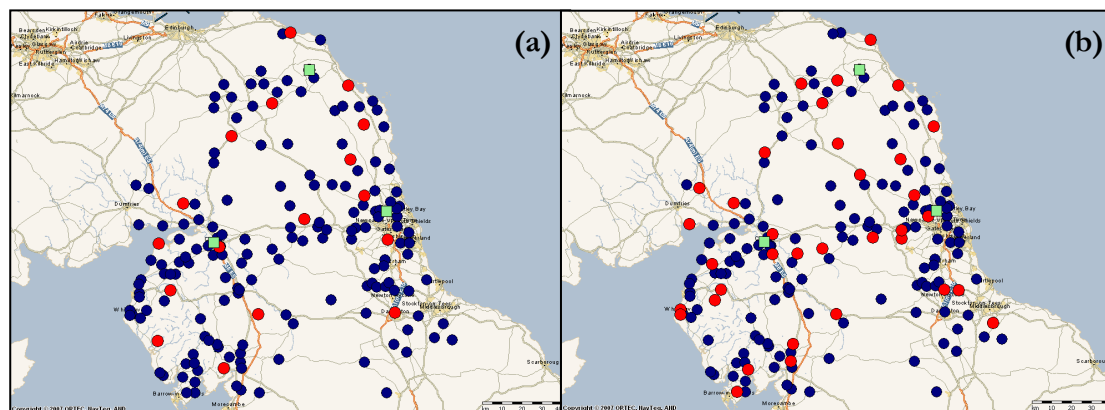


Figure U.2. The seeds are given in red, the depots are green, and all the orders are blue. (a) represents the seed selection procedure of the Period Scheduler, balanced on one visit per depot per day, (b) represents the flexible clustering method.

Clustering method	Num. of orders	Volume (litres)	Dist. (km)	Volume/kilometre (litres/km)	Cost (fictive)	Cost/litre	Veh. used	Average relative delivery	Volume/kilometre improve.	Cost/volume savings
Base scenario*	648	718.457	13.448	53,4248	41.761	0,0581	56	53,55%	-	-
Flexible clustering	978	981.942	15.095	65,0508	49.744	0,0507	68	59,77%	21,76%	12,85%
PS cluster 18 seeds	922	956.837	15.191	62,9871	50.072	0,0523	72	59,55%	17,90%	9,97%
PS cluster 36 seeds	976	988.631	16.538	59,7794	53.198	0,0538	69	59,94%	11,89%	7,43%

* The base scenario was planned by the planners, and therefore has no clustering method.

Table U.5. The results illustrate that the method for clustering used in the Period Scheduler does not lead to more savings in volume per kilometre, but is close.

Additionally, the lower number of seeds (18 seeds when there is one seed per depot per day) makes the customer assignment phase a lot faster. The time for this phase reduces from around 1.000 seconds (~18 minutes), to 260 seconds (~4,3 minutes). When we use 36 seeds, the algorithm run time is the same as in the flexible clustering case.

A disadvantage of the Period Scheduler clustering method with 18 seeds is that all vehicles go to one area on each day. One of the solution requirements is that the vehicle should go to multiple areas on one day, to cope with emergency orders if they exist.



Figure U.3. The geographical spread for week 42 with the clustering method of the Period Scheduler and 18 seeds illustrate that Tuesday has one depot without workload.