Master thesis, Industrial Engineering and Management

Improving the waste collection planning of Amsterdam

Author

W.J.P Heijnen

w.j.p.heijnen@student.utwente.nl

Committee of supervisors

Dr.ir. M.R.K. Mes

Dr.ir. E.A. Lalla

University of Twente, dep. Industrial Engineering and Business Information Systems

H. Palm

DAT.Mobility, CS Conculting

Ing. A.H. Oldenburger

Goudappel Coffeng, MD Mobiliteitsdiensten

24-06-2019

Management summary

Introduction

This research is conducted at the Goudappel Group, where it is part of a project initiated by the municipality of Amsterdam. Amsterdam's motivation for this project is to improve the performance of its waste collection planning. The waste collection of Amsterdam involves the collection of the waste generated by its 854 thousand inhabitants. The scope of this research is limited to the waste collected from underground containers, of which there are 12.000 scattered over Amsterdam.

The goals of the municipality of Amsterdam are to reduce the total collection costs, distance driven, and improve the service level to its inhabitants. The current planning methodology, which is based on fixed, static schedules and emptying frequencies, is unsuitable for the stochastic and dynamic demand for waste collection. Accordingly, the main research question of this research is: in which way and to what degree can the waste collection planning of Amsterdam be improved by using dynamic scheduling algorithms?

The main research question is answered by analyzing the characteristics of the waste collection system of Amsterdam, reviewing existing literature to find related problems, proposing a new solution approach based on the findings of the literature review, and evaluating its performance on the waste collection planning of Amsterdam to formulate recommendations for the municipality.

Waste collection in Amsterdam

The logistical chain that is set up in Amsterdam to collect its waste consists of four types of locations: (1) containers, used to collect waste, (2) wharfs, which function as a base for collection operations, and two types of disposal facilities where collection vehicles can dispose the waste they collected: (3) satellite facilities and (4) waste processing facilities.

During the planning of the waste collection, decision makers should make two decisions: when to empty which container and how to combine these containers into vehicle routes. This problem most closely resembles the theoretical Inventory Routing Problem (IRP).

Proposed solution approach

To solve this problem, this thesis proposes a novel solution approach that consists of three phases that are executed on a rolling horizon:

1. Container selection, where containers that are expected to be relevant within the pre-determined planning horizon are selected.
2. Day assignment, where the selected containers are assigned to days of the planning horizon.
3. Route construction, where routes are constructed to collect all assigned containers.

The main focus of this research on the second phase, where a novel method is proposed to consider both the time and space dimensions of the IRP. These dimensions are important to consider simultaneously, as decisions in one dimension influence and restrict the possible decisions in the other dimension.

During the day assignment, it is decided which containers are emptied on which day. During this decision, both the timing and location of each container is considered. To do this, two new costs approximation measures are introduced: the cluster fitness approximation, related to the distance to other containers scheduled on the same day, and the timing penalty costs, related to the costs of emptying a container earlier or later than its expected optimal emptying day. The container is assigned to the day for which the sum of the two approximations is lowest.

Experiment results

To evaluate the performance of the proposed solution approach, a simulation model is implemented that models the Zuidoost-district of Amsterdam. During the numerical experiments, the proposed solution approach is adjusted to suit the characteristics of Amsterdam Zuidoost using several experimental factors that influence planning decisions taken in the proposed solution approach. The numerical experiments show that considerable improvements are possible if a dynamic planning approach is adopted. Moreover, the experiments show the potential benefits of installing fill level sensors in containers. Figure 1 shows the possible range of performance without and with installing sensors, compared to the performance of the current planning methodology of Amsterdam. The results show that a travel distance reduction of 12% is possible without installing sensors, and 21% with sensors, without reducing the service level offered to Amsterdam's inhabitants.

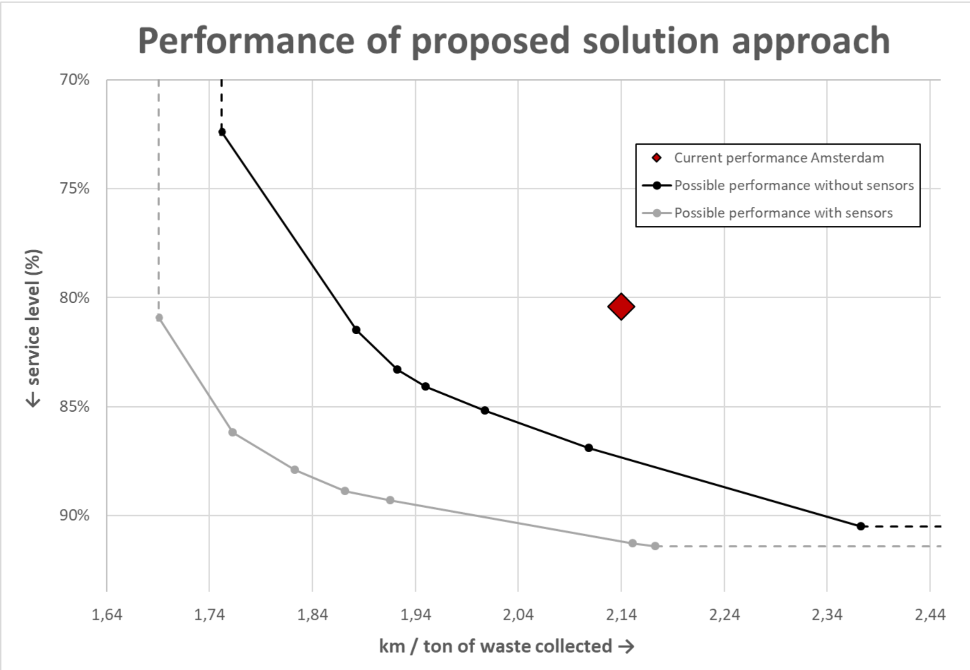

Figure 1: Possible performance of the proposed solution approach compared to Amsterdam

Recommendations

Based on the results of the experiments we formulate four recommendations for the municipality of Amsterdam:

- Implement the proposed solution approach using one of the preferred configurations,
- Implement sensors into all or part of all containers,
- Improve the quality of data collection to enable better fill level predictions,
- Initiate further research into the benefits of adding additional dynamicity.

Discussion and further research

We believe that this thesis introduces a novel solution approach to the IRP in the way the time and space dimensions are considered simultaneously. Moreover, the solution is generalizable and can be applied to waste collection problems of other cities and probably even other IRPs. Further research should be done to confirm these presumptions. Moreover, the cost approximation methods, i.e., the cluster fitness approximation method and timing penalty costs, should be subject to more research as these approximations are currently merely very rough approximations of the costs associated with decisions taken in the IRP. It is believed that more accurate approximations of these costs can considerably improve the solution quality of the proposed solution approach.

Preface

The thesis before you is the result of my Master's assignment, which I conducted to receive my degree in Industrial Engineering and Management at the University of Twente. My assignment, as offered by DAT.Mobility and Goudappel Coffeng, was to improve the current planning methodology used by the municipality of Amsterdam for their waste collection. This proved to be quite a challenge, during which I have learned a lot.

I would like to thank André Oldenburger for providing me with the opportunity to work on this project and for actively involving me in the 'project team' to gain some real insight and experience of working on such large projects. Furthermore, I want to thank Henri Palm for all his feedback, which was always positive, but above all constructive. It was a pleasure working at DAT.Mobility where there was always a positive and open atmosphere.

During this project, I was also advised by my supervisors from the University of Twente, Martijn Mes and Eduardo Lalla, whom I would also like to thank. Their critical feedback on the previous versions has undoubtedly greatly improved the quality and clarity of this thesis.

Wouter Heijnen

Enschede, June 2019

Table of contents

List of definitions

To prevent ambiguity we start by listing some definitions that are used throughout this thesis. Subsequently, a list of all abbreviations is given.

Tour	A tour is a sequential list of containers that are subsequently visited that starts at a wharf or disposal facility and ends at a disposal facility.
Route:	A route is a sequential list of tours that are subsequently executed by a vehicle. A route consists of one or multiple tours.
Scheduling/schedule	Scheduling encompasses the operational decisions that are taken to make the schedule. A schedule is a collection of routes, which vehicles should perform which routes, and when.
Planning:	Planning denotes decisions of a higher hierarchical level than scheduling and is used to comprise all decisions taken by the waste management department to plan the waste collection.
Container/vehicle fill level	The fill level of both containers and vehicles indicates the amount of waste that is deposited and is currently stored in said container or vehicle. The fill level is often displayed as a percentage of the total capacity of the container or vehicle.
Container deposit rate:	The container deposit rate indicates the speed at which waste is deposited in a container. Because the deposit rate is stochastic, it is represented as a statistical distribution.

Chapter 1 - Introduction

This chapter introduces the research and the problem that initiates it. Firstly, we discuss relevant background information about the affiliated organizations and their connections to this research (Section 1.1). Secondly, the problem description is given, outlining the core problem this research aims to solve (Section 1.2). Thirdly, the research design is discussed in two sections: the research objectives and the research questions (respectively in Section 1.3 and 1.4). Lastly, the problem solving approach and thesis structure is outlined (Section 1.5).

1.1 - Research context

This research takes place at the Goudappel Group, a collaboration of companies that work together to offer their customers integral solutions to mobility issues. The offices of Goudappel Group are located in several cities across the Netherlands. This research is conducted from the headquarters, located in Deventer. The two most prominent parts of the Goudappel Group are Goudappel Coffeng and DAT.Mobility. The largest, Goudappel Coffeng, is mainly focused on consulting in mobility issues. Whereas DAT.Mobility supports Goudappel Coffeng, as well as its own customers, with data analysis and the development of ICT solutions. Both companies take part in this research. In the remainder of this report, the Goudappel Group is referred to as Goudappel.

Recently, Goudappel started to expand, from their traditional position as advisors in mobility, into the field of logistics. The ambition is to combine their pre-existing knowledge and experience in traffic modeling with a transport planning application. The desired end result is a dynamic planning platform that is usable in a broad range of logistic applications and scenarios. This research is the first step towards developing this dynamic model. The first practical application to present itself to Goudappel is the waste collection planning of Amsterdam. Therefore, while this theoretical research is conducted at Goudappel, it is part of a practical project initiated by the municipality of Amsterdam to improve its waste collection planning.

Amsterdam is the capital and the largest city of the Netherlands and has more than 854 thousand inhabitants (Municipality of Amsterdam, 2018b). The municipality of Amsterdam is responsible for the collection of all the waste produced in Amsterdam. Waste collection is one of the most complex and visible services offered by municipalities and involves large expenditures. In recent years, municipalities are therefore increasingly reconsidering their waste management due to costs and environmental concerns (Nuortio et al., 2006) (Jewel, 2017), as is now the case in Amsterdam. The problem of planning the collection of waste in Amsterdam is particularly complex because of the large number of densely packed containers and narrow, congested streets. Moreover, the addition of multiple satellite facilities, waste processors, and wharfs, which are elaborated upon in Section 2.2, complicates the planning problem.

1.2 - Problem description

Amsterdam's main motivation for initiating the project is its ambition to improve the performance of its waste collection planning. The main targets of Amsterdam are to reduce the total cost of collection and to minimize the number of required vehicle movements. Concretely, the municipality wants to reduce the cost of waste collection from €43 million to €40 million per year (Municipality of Amsterdam, 2018c). The targets are motivated by three performance indicators on which the municipality of Amsterdam perceives itself to be underperforming: the costs of waste collection, the service to its inhabitants, and the emission of greenhouse gases. These targets are set for the entirety of Amsterdam's private, so non-industrial, waste collection, consisting of: collection from waste containers, bulk collection, and manual collection (Municipality of Amsterdam, 2018c). However, this research chooses to focus solely on the collection from waste containers, as this collection method is believed to have the highest potential for improvement.

To find the core problem that causes the underperformance of the waste collection planning, a problem cluster is drawn, as described by Heerkens and van Winden (2012). This is used to analyze the surrounding problems and deduce their causal relationships. In the problem cluster, which can be found in Appendix 1, two main themes can be identified: the non-optimal timing of emptying containers and the inefficient collection routes.

The first theme, the non-optimal emptying time of containers refers to containers being emptied either too early or too late. The main reason for this is the static nature of the planning methods used by Amsterdam. Containers are emptied with predetermined frequencies that do not consider an up-to-date forecast. Because the schedule is made to effectively collect waste under all circumstances, the schedule and corresponding container emptying frequencies are determined to be as robust as possible. However, this causes inefficiencies as the content of containers is stochastic, changes over time, and is subject to the effects of seasonality. The combination of these factors and the static planning it can, for example, occur that a container is emptied when relatively empty because it is a less busy period of the year, or containers overflow because of an increase in residents in the neighborhood. Both these situations are undesirable as overfull containers inconvenience inhabitants and signify a bad service level, while containers that are emptied too early cause unnecessary vehicles kilometers and emissions.

The second theme, the inefficiency of the routes between containers can be attributed to two factors: the decentral planning and a failure to respond to new incoming information, such as traffic data or calling inhabitants. Firstly, Amsterdam's waste collection planning is decentralized, this means that each municipal district has the responsibility for its own waste collection. Consequently, containers are designated to the district in which they are located and planned in routes executed by that district, even though it could fit more efficiently in a route of another district. This causes detours and unnecessary additional driven kilometers. Secondly, the current routes do not consider predicted or real-time traffic data. Especially in densely populated areas such as Amsterdam, where traffic congestion is a recurring phenomenon, this data should be used to avoid such congestions. The inability to avoid traffic congestion threatens the feasibility of the schedule and causes additional CO_2 emissions.

To summarize, the two main problem themes are caused by, in the first case, static scheduling, and in the second case, static routing. This static scheduling and routing policy, in combination with the changing waste collection setting in Amsterdam, causes the underperformance of the planning. As a result, the core problem is formulated as follows:

Core problem: *The current static collection schedules and routes are unsuitable for the stochastic, dynamic demand for waste collection.*

1.3 - Research objective

The aim of this research is to solve the core problem as described in Section 1.2. Both Goudappel and Amsterdam agree that the development of a new dynamic planning platform for the collection of waste is the most suitable solution for this problem. Utilizing dynamic planning techniques allows individual containers to be scheduled separately based on the forecasts of their fill levels. This gives the opportunity to improve the timing of emptying containers. Moreover, the new planning techniques can incorporate improvements on the other previous weakness as well, such as the inefficient routes and not considering traffic data.

Because developing an entire planning platform is unrealistic in the time allotted for this research, the scope is limited to the collection of waste from waste containers. This collection method is chosen as it is highly influential for the total waste collection performance, easily separable from the other planning functions, and not subject to as many legislative restrictions. Moreover, the forecasting of container content is such a complex problem that it is considered largely outside the scope of this research. It is briefly discussed as an input variable for the planning, but a complete, thorough analysis of all the complexities involved is not conducted. Lastly, developing an interface and implementing the planning application are also omitted from the research scope.

At the moment, neither Goudappel nor Amsterdam have the necessary experience or knowledge to develop the required planning methods. Therefore, the goal of this research is to identify and develop a planning methodology that improves upon the current waste collection planning in Amsterdam. More specifically, the new planning method should improve upon the current cost, emission, and service performance indicators.

1.4 - Research questions

From the previously described research objective we deduce the following research question:

Main research question:
In which way and to what degree can the waste collection planning of Amsterdam be improved by using dynamic scheduling algorithms?

To answer the main research question, we formulate several research questions. These research questions are formulated below and are elaborated upon further in Section 1.5.

Research question 1: What is the current state of the waste collection planning in Amsterdam?

a. What are the characteristics of the waste collection system in Amsterdam?
b. How is the collection of waste currently planned?
c. What are the relevant performance indicators to evaluate the waste collection planning?
d. What is the performance of the current waste collection planning in Amsterdam?

Research question 2: What relevant routing problems and corresponding solution approaches are described in the literature?

a. What types of existing routing problems most closely resemble the problem faced during the waste collection planning in Amsterdam?
b. What types of solution approaches are described in literature to solve these types of problems?
c. What techniques are applied in the literature to improve waste collection planning?
d. How can the performance of different planning methodologies be tested and evaluated?
e. How are the findings of this literature review applicable to the context of Amsterdam?

Research question 3: How can a novel planning methodology be designed for the waste collection in Amsterdam?

a. Which decisions are to be taken to schedule the collection of waste in Amsterdam and which boundary conditions apply to those decisions?
b. How can the decision process be decomposed into smaller decisions?

Research question 4: How should the waste collection system of Amsterdam be modeled to allow for the evaluation of novel planning methodologies?

a. In what scope and level of detail should the planning methodologies be tested?
b. How should the appropriate values for the adjustable parameters of the proposed solution approach be established?

Research question 5: What is the expected outcome of the proposed planning methodology for the waste collection of Amsterdam?

a. What are the effects of the defined adjustable parameters on the planning performance?
b. How does the proposed solution approach deal with changing demand for waste collection?

1.5 - Problem solving approach

To solve the core problem and achieve the research objective, the research questions formulated in the previous section should be answered. This section outlines the problem solving approach that is adopted during this research to answer the main research question and in doing so gives an overview of the structure of this thesis.

The first research question is answered in Chapter 2, which gives an overview of the current situation of Amsterdam's waste collection system and its planning. To give this overview, the characteristics of the waste management system in Amsterdam are discussed, as well as the current planning methodology. Subsequently, relevant performance indicators are chosen with which to evaluate the current planning performance using data supplied by the municipality of Amsterdam.

The second research question involves a review of the current literature relevant to the problem in Amsterdam. This literature review is given in Chapter 3 where we discuss literature on different routing problems, solution approaches, and waste collection planning. The applicability of the discussed literature to the context of Amsterdam is debated and finally, the contribution of this research to the current literature is discussed.

Based on the previous literature review, the third research question asks how a new planning methodology can be designed for the waste collection of Amsterdam. Chapter 4 starts by defining the problem faced by Amsterdam and continues by formulating our proposed solution approach. Several tunable experimental parameters are identified and introduced.

The fourth research question discusses how the proposed solution approach can be compared to the current planning methodology of Amsterdam. Moreover, how the effects of the experimental parameters can be evaluated. To answer this question, Chapter 5 describes the chosen modeling approach and the implemented model, including the chosen scope and level of detail.

In the fifth research question, we study the expected performance of the proposed solution approach and the effects of the tunable parameters using numerical experiments. These experiments are discussed in Chapter 6. Chapter 7 formulates the conclusions based on the results of the experiments and gives recommendations for the municipality of Amsterdam. Moreover, the research is discussed and recommendations are given for further research.

Chapter 2 - Current state of waste collection in Amsterdam

This chapter describes the current state of the waste collection in Amsterdam in order to clarify the context in which the previously described problems occur. It starts by giving an introduction to some general characteristics of waste collection in Amsterdam (Section 2.1). After that, the logistical chain set up to handle the collection of waste (Section 2.2) and the available collection vehicle fleet (Section 2.3) are discussed. Following that, the waste collection process is discussed briefly (Section 2.4), as well as the current planning methodologies applied by Amsterdam (Section 2.5). This chapter continues by formulating the relevant performance indicators (Section 2.6) and evaluating the performance of the current planning methodology (Section 2.7). Finally, the research questions pertaining to this chapter are answered in a summarizing conclusion of this chapter (Section 2.8).

2.1 - Waste collection in Amsterdam

The waste collection of Amsterdam is the responsibility of its municipality. Annually, it facilitates the collection of more than 304 thousand tons of waste, which is done by around 300 full-time employees (Municipality of Amsterdam, 2015). In accordance with the scope of this research, we focus on the collection of waste from the more than 12.000 waste containers scattered in Amsterdam, see Figure 2.

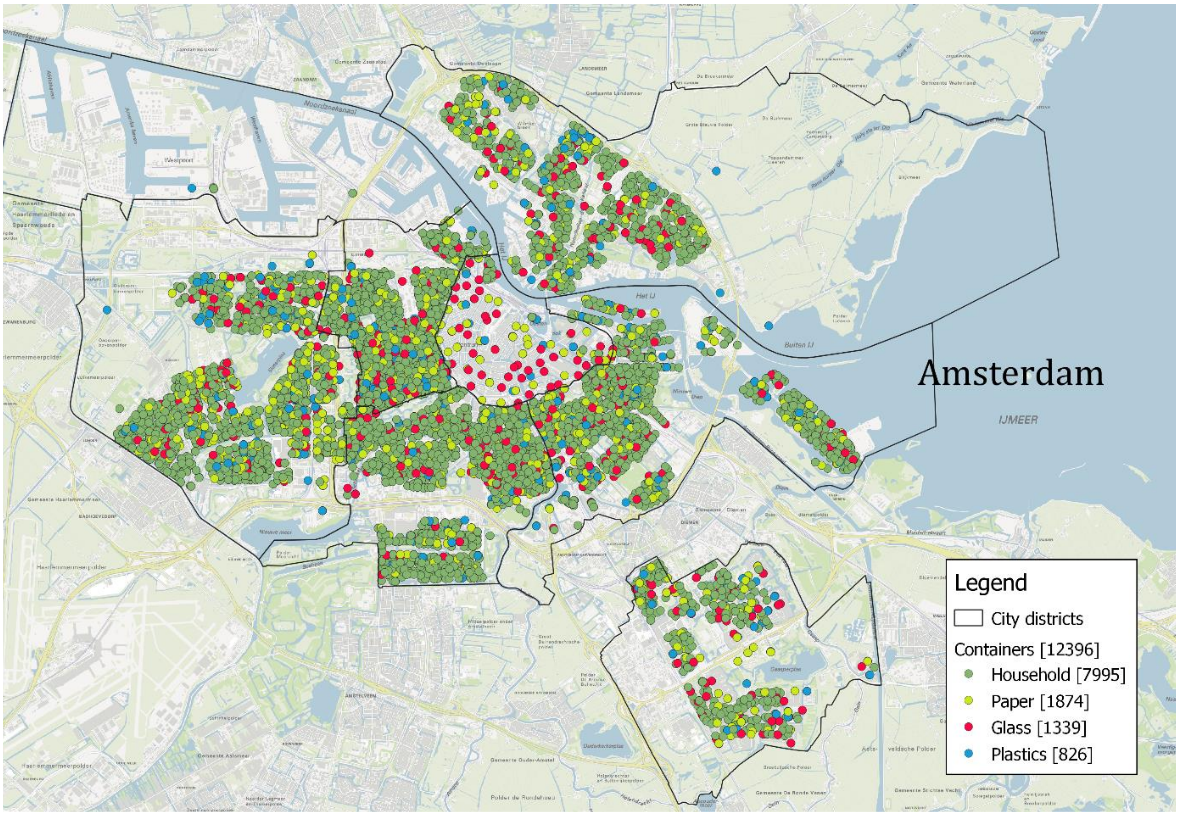

Figure 2: Dispersion of waste containers per fraction over Amsterdam's districts

Amsterdam distinguishes six different types of waste, called waste fractions, that are collected and processed separately. Four of these waste fractions are wholly or partly collected in waste containers: household waste, glass, paper, and plastics. Household waste can also be described as residual or unsorted waste. On average, each inhabitant of Amsterdam dumps a total of 227 kilograms of waste in waste containers per year. 81% of that waste is thrown out as unsorted household waste. The remaining 19% is sorted into either the designated paper, glass, or plastics containers (Municipality of Amsterdam, 2015). Appendix 2 shows the distribution of the weight collected from containers per waste fraction, per year.

2.2 - Logistical chain

The logistical chain set up by Amsterdam to collect waste consists of four types of locations: waste containers, wharfs, and two types of disposal facilities: satellite facilities and waste processing facilities. Each location has its own functions (Municipality of Amsterdam, 2018c) and characteristics which are described in the following sections.

2.2.1 - Waste containers

Waste containers are used to collect and store waste. Inhabitants can throw away waste in the waste containers, in which it is then stored out-of-sight and without smell, improving the cityscape. The storage of containers can be located underground, as in the example in Figure 3, or above ground. Waste is collected from the containers, to ensure they do not overflow, by waste collection vehicles. Each container is dedicated to one waste fraction to enable the municipality to recycle the separate waste fractions. Containers have a limited waste storage capacity which is between $3m^3$ and $7m^3$.

Figure 3: Emptying an underground waste container (photo credit: Palfinger.ag)

2.2.2 - Wharfs

Wharfs are used as the base for all collection operations in a city district. Waste collection employees often have fixed designated wharfs where they start and end their working day. All waste collection routes also start and finish at a wharf. The wharfs then function as parking lots for all collection vehicles overnight. Because of the current decentralized way of planning, each of the seven districts has a wharf from which all waste of that district is collected.

2.2.3 - Waste disposal facilities

There are two types of locations where collection vehicles are able to dump their collected waste: satellite facilities and waste processing facilities. The difference between the two locations is that the satellite facility is only used to temporarily store waste, which is later transported to a waste processor to be processed. The waste processors are the end of the logistical chain for all collected waste in

Amsterdam, where the waste is separated and recycled or processed if possible. Whether to visit a satellite facility or waste processor is a planning decision. This decision is based on the proximity of the route to either of the disposal facilities.

There are currently two satellite facilities with plans to build a third. The two existing satellite facilities have a combined capacity to store the waste load of 55 full collection vehicles. Moreover, there are four waste processing facilities, each of which is dedicated to processing waste from one waste fraction.

2.2.4 - Geographical layout

The locations of the wharfs, satellite facilities, and waste processing facilities are shown in Figure 4. Note that the satellite facility planned for construction is already included in this figure, see the northern-most satellite facility, as its construction starts in 2019.

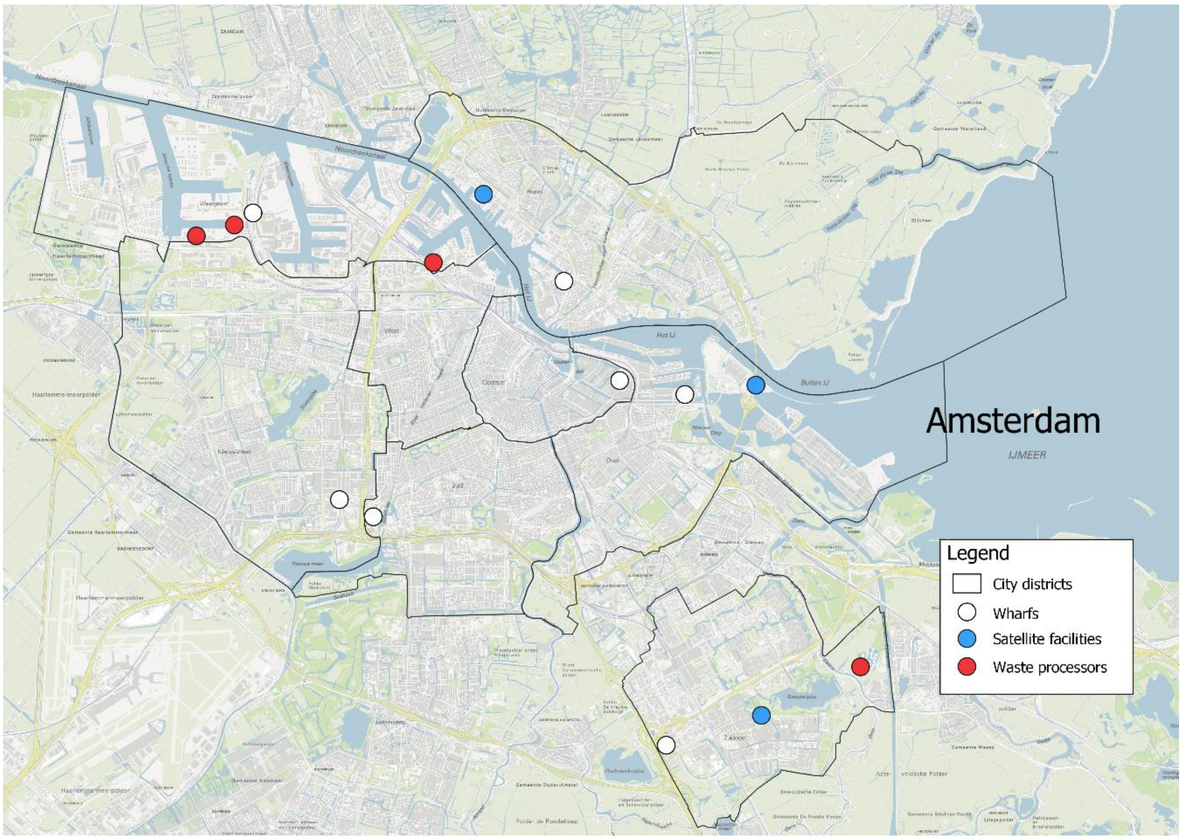

Figure 2: Locations of wharfs, satellite facilities, and waste processors in Amsterdam

2.3 - Available collection vehicles

This section describes the collection vehicles that are available for waste collection in Amsterdam. Waste collection vehicles have several characteristics that dictate how they can be utilized. The most notable are their compatible waste fraction, capacity, and the capability to unload waste at satellite facilities.

The municipality of Amsterdam has a total of 122 waste collection vehicles at its disposal. More than half of these, 62, are used to collect waste from waste containers. The remaining vehicles are used for bulk and manual waste collection. Each vehicle is currently transfixed to one designated wharf, but this is not necessarily the case in the future.

Collection vehicles can be categorized as either dedicated or multipurpose vehicles. Dedicated collection vehicles are compatible with only one waste fraction, while multipurpose vehicles can be utilized to collect multiple different waste fractions. Multipurpose vehicles are not able to collect different waste fractions simultaneously, but only consecutively. To illustrate, a multipurpose collection vehicle can collect paper in the morning and household waste in the afternoon, as long as it is not on the same trip. Table 1 shows the number of available dedicated and multipurpose vehicles per waste fraction (Municipality of Amsterdam, 2018c). Note that multipurpose vehicles are counted for all waste fraction with which they are compatible.

Table 1: Available waste collection vehicles

Waste fraction	Household waste	Glass	Paper	Plastics	Total
Dedicated	35	1	2	4	**42**
Multipurpose	12	18	20	0	**20**
Total available for fraction	47	19	22	4	**62**

Furthermore, collection vehicles have a limited capacity to store and transport waste. Most collection vehicles have a weight capacity between 8 and 10 tons of waste. Several collection vehicles are also equipped with a storage tank-switch system. This system enables collection vehicles to switch their full storage tank with an empty tank at one of the satellite facilities. Collection vehicles can then resume collection waste without visiting a waste processor.

2.4 - Waste collection processes

To gain a better understanding of the waste collection in Amsterdam this section describes the associated processes. This is done using business process modeling as it facilitates the understanding and analyzing of processes (Aguilar-Savén, 2004). The process flow chart is shown in Figure 5 (Municipality of Amsterdam, 2018c).

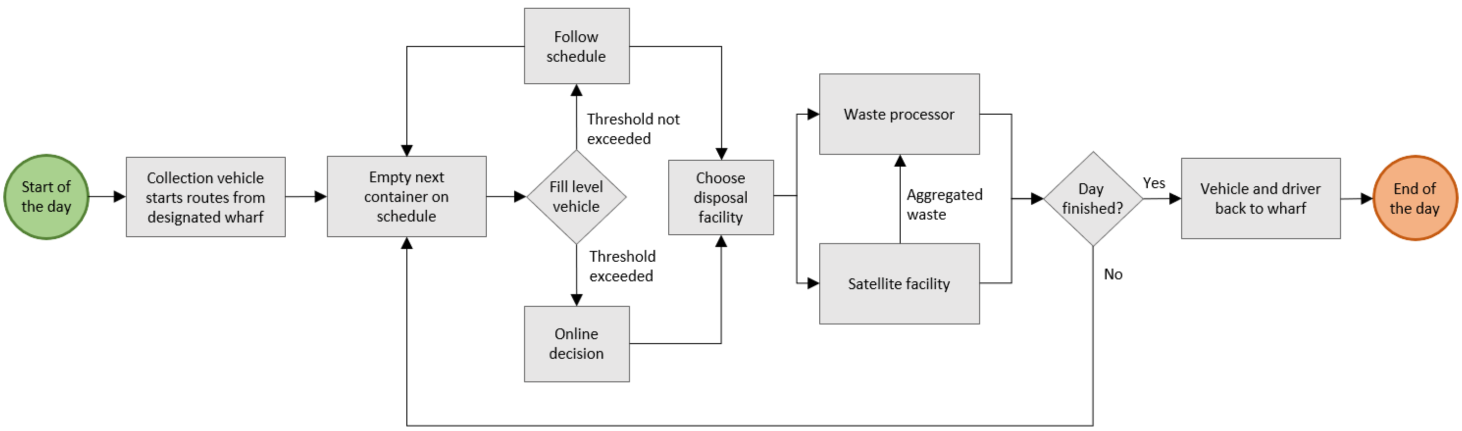

Figure 3: Process flow chart of a day of one waste collection vehicle in Amsterdam

At the start of each working day, which are Monday to Friday, the waste collection routes start from the different wharfs. Each employee ordinarily works from the same wharf every day. The wharfs also serve as an overnight storage location for waste collection vehicles. A typical working day consists of 8.5 hours, of which, after subtracting time for activities such as preparation and lunch, 6.25 hours are available to execute the planned collection routes.

During the execution of the collection routes, there is a decision moment every time a container is emptied. During this moment, it is considered if it is possible for the collection vehicle to continue according to the predetermined schedule. This decision is based on the fill level of the collection vehicle, if a capacity threshold is reached, the vehicle is considered too full to continue to empty another container. If that is the case, an online decision is taken to either dump the collection vehicles' load at a satellite facility or waste processor. This online decision is necessary because the actual fill level of containers is uncertain and they may have higher fill levels than anticipated, resulting in infeasible routes.

During the execution of the collection routes, there is a decision moment after each emptied container. The decision is based on the fill level of the vehicle; if the vehicle is too full to empty another container it should decide to empty its storage at either a waste processor or satellite facility. If the vehicle is not too full, the normal schedule can be followed which is either: collect waste from the next container or go to either the waste processor or satellite facility. A waste collection vehicle can, for example, have a higher fill level than anticipated because of higher fill levels of the collected waste containers. When an employee is finished with his/her collection route, the collection vehicle can be retired to the wharf and the working day is finished. Occasionally, the waste that is aggregated at the satellite facilities is transported to the corresponding waste processor.

2.5 - Current planning

The current schedules are static and fixed, this means that the same schedules are executed in a cyclic manner. Moreover, the routes between waste containers in these schedules are also rarely subject to changes (Municipality of Amsterdam, 2018c). The schedules are currently based on a predetermined emptying frequency per container, e.g., once per week, twice per week, once per two weeks. These frequencies are based on slightly exaggerated estimations of container deposit rates to ensure the frequencies can be used year-round.

The static nature of the schedules means that planners are mostly occupied with the operational issues, such as adjusting the planning when containers are defect or when inhabitants complain about uncollected waste. Inhabitants are able to notify the municipality using a public notification portal. Amsterdam strives to handle all complaints within three days. In reality, this is managed for 80% of the complaints (Municipality of Amsterdam, 2016). Appendix 3 shows the five most common complaints and notifications concerning waste collection the municipality of Amsterdam received in 2014 (Municipality of Amsterdam, 2015). When a complaint is made, an addition has to be made to the collection planning, this is currently done manually by Amsterdam's planners. Next to the operational planning of the schedules, the planner is also responsible for the personnel and vehicle planning.

Currently, district planners are increasingly trying to work together to improve the existing schedules. However, they are severely limited by the applications they have access to. Amsterdam does not have an integrated planning application, moreover, the districts do not have the same applications. Therefore, when they want to collaborate, they have to communicate their schedules and routes using Excel, Word, or physical maps.

2.6 - Choosing performance indicators

This section starts by determining the key performance indicators (KPIs) on which the waste collection planning should be evaluated. These KPIs are chosen based on a combination of wishes from the municipality and suggestions from literature. Because multiple KPIs are identified, we also discuss how alternatives can be evaluated based on multiple, possibly conflicting, KPIs.

2.6.1 - Choosing key performance indicators

The three main targets for Amsterdam in their pursuit for an improved waste collection planning methodology are: to reduce the total costs and vehicle movements and to improve the service to its inhabitants (Municipality of Amsterdam, 2018c). In addition to these main KPIs, the municipality also formulates several secondary objectives: total CO_2 emission, time spent collecting, average fill level of containers upon emptying, and fill level collection vehicles upon unloading. The three main KPIs are defined and quantified as follows:

- Influenceable collection costs / ton of waste collected: consisting the sum of the vehicle, fuel, and wage costs;
- Number of kilometers driven / ton of waste collected;
- Service level: the percentage of waste that is collected on-time.

2.6.2 - Multiple criteria decision making

Evaluating alternatives given a set of criteria, or KPIs as named in this research, is defined by Stewart (1992) as Multiple Criteria Decision Making (MCDM). The aim of MCDM is to assist the decision maker to find the most desirable alternative and provide justification for that decision (Stewart, 1992).

To evaluate different planning methodologies, we use the concept of Pareto optimality. Pareto optimality is used to evaluate alternatives in multiple papers on multi-objective general routing (Huang, Fery, Xue, & Wang, 2008), as well as waste collection routing (Xue & Cao, 2016) (Samanlioglu, 2013). Due to the often conflicting natures of the objectives, oftentimes, no one solution is objectively better than all other solutions. Alternatives are called Pareto optimal if an improvement in one objective has to be at the expense of another objective (Huang et al., 2008). The set of all Pareto optimal points is called the Pareto curve or front, an example is shown in Figure 6. This front gives insight into the shape of the trade-off between the objectives and can be used during the decision making.

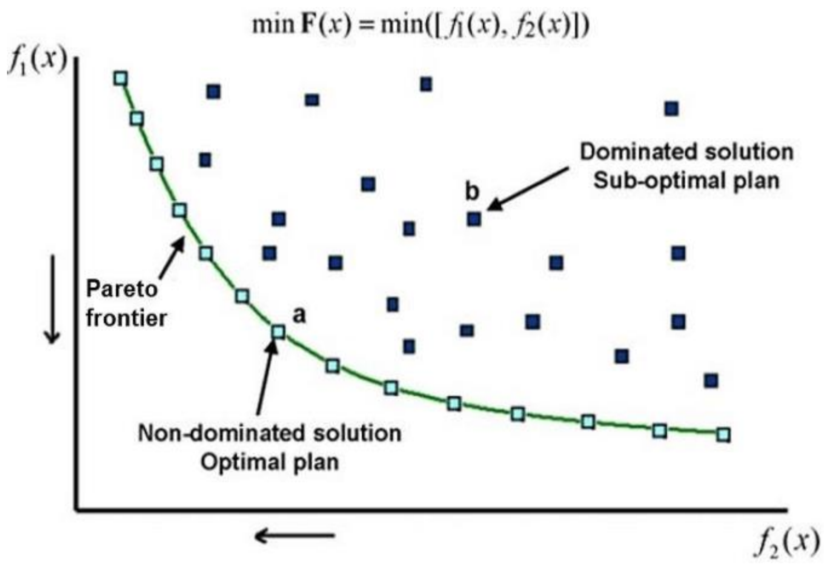

Figure 4: Example of Pareto frontier (Di Somma, 2016)

To allow for a straightforward evaluation of alternatives and to reduce the number of Pareto optimal point to consider, we choose to select two criteria on which to evaluate all alternatives. More criteria would complicate the analysis and selection of promising alternatives, as two-dimensional visualization is then made impossible. The two KPIs used for the multi-objective decision making are the number of kilometers driven per ton of waste collected and the service level. These two criteria represent both sides of the most important trade-off: emptying early and often versus emptying late and infrequent. In the remainder of this thesis, the Pareto front is called the efficiency frontier.

2.7 - Current performance

This section discusses the performance of the current waste collection planning on both the primary and secondary KPIs identified by Amsterdam, with the exception of the service level. This KPI is excluded because the municipality currently does not have the ability to accurately measure it. Moreover, the municipality of Amsterdam has only gathered the performance of the planning of the household, paper, and glass waste fractions. Because the plastic fraction represents such a marginal part of the total collected waste, i.e., less than 1% (Municipality of Amsterdam, 2015), we presume that the other fractions give an accurate depiction of the current planning performance. The performance data relates to the waste container collection performance of 2017.

2.7.1 - Influenceable collection costs

The influenceable collection costs consist of four factors: the cost per vehicle, personnel costs, cost per kilometer, and satellite facility costs. The contribution of each cost factor is shown in Table 2 (Municipality of Amsterdam, 2018a).

Table 2: Influenceable costs incurred by Amsterdam in 2017 for the collection of household, paper, and glass from containers

Cost factor	Amount		Costs	% of costs
Required number of collection vehicles	67	€	6.375.000	50,1%
Hours spent collecting	137.152	€	5.486.090	43,1%
Number of kilometres driven	940.735	€	658.514	5,2 %
Containers via satellite facilities	8.256	€	643.930	5,1 %
Influenceable collection costs	**-**	**€**	**12.721.507**	**100 %**

The highest incurred costs are the fixed costs of the collection vehicles. On average maintaining and depreciating the vehicle fleet costs approximately €95.000 per vehicle, per year. The personnel cost, or the time spent collecting, is a close second in terms of impact on the total costs. In 2017, collection employees worked a total of 137.152 hours for an average cost to the municipality of €40. The latter two have a significantly less pronounced impact on the total costs, cumulatively around 10%, and are calculated using fuel costs of €0,70/km and satellite facility costs of €78/container.

While most costs are incurred by the household waste fraction, 77%, this is also the most prevalent waste fraction. To fairly compare the cost-effectiveness of the waste fractions, we should take the amount of collected waste into account. This comparison is shown in Table 3, where the costs are shown from two perspectives: per ton of waste collected and volume of waste collected. This distinction is relevant because of the difference in density of different waste fractions. The densities are estimated as follows: household (100 kg/m^3), paper (70 kg/m^3), and glass (300 kg/m^3) (Municipality of Amsterdam, 2018a). Table 3 shows that the collection of household waste is most cost-effective, as the cost per ton and per cubic meter collected are lower than average. At the same, presumably caused by the respective densities of the waste fractions, the collecting cost per ton of paper and per cubic meter of glass are significantly higher than average.

Table 3: Costs per ton and cubic meter of waste collected

Waste fraction	Household	Paper	Glass	Average
Cost per ton of waste	€ 55,15	€ 120,29	€ 55,79	**€ 60,79**
Cost per m^3 of waste	€ 5,52	€ 8,42	€ 16,74	**€ 6,14**

2.7.2 - Average fill levels of containers and collection vehicles

The fill level is the percentage of the total volume capacity that is filled with waste. The moments at which the container is emptied and the collection vehicle dumps its waste at a disposal facility are used to measure the average fill levels. Table 4 shows the fill levels of both containers and vehicles for the collection of different waste fractions, as well as the average fill levels.

Table 4: Average fill levels of containers and vehicles upon emptying

Waste fraction	Household	Paper	Glass	Average
Average fill level of containers	39,3 %	39,6 %	20,3 %	**38,1 %**
Average fill level of vehicles	79,1 %	74,0 %	75,4 %	**77,9 %**

It can be seen that the average fill level of containers is especially low, as previously discussed in Chapter 1. Notably, the average fill level of glass containers is almost half that of the other waste fractions. This is caused by the fact that the fill level of containers is calculated based on weight. As the volume capacities for the average container are comparable, the high density of glass means that glass containers have a very high weight capacity. Therefore, while paper containers are emptied twice as much with a comparable amount of waste, the fill level of glass containers is approximately twice as low.

2.8 - Conclusion

This chapter gives an overview of the current state of waste collection and its planning in Amsterdam. In doing so, it answers research question 1: What is the current state of waste collection planning in Amsterdam?

A distinctive characteristic of the waste collection logistical chain in Amsterdam is the presence of satellite facilities. Satellite facilities can be used to offload full collection vehicles, allowing them to collect more waste. The current collection planning is predominantly fixed and cyclic. Changes in the planning that do occur are caused by inhabitant complaints or defective equipment. To evaluate the current planning and to be able to compare it to new proposed planning methodologies, we formulate four KPIs by consulting the municipality and literature: influenceable collection costs, number of overfull containers, and the average fill levels of both containers and vehicles. Notably, by far the most costs are spent on the fixed costs of vehicles (47%) and the hourly wages of collectors (43%). Another noteworthy finding is the low average fill level at which containers are emptied (38%).

The next chapter outlines a literature review which places Amsterdam's problem in the existing literary theory and shows how similar problems are approached and solved in previous literature.

Chapter 3 - Literature review

This chapter reviews the relevant literature to solve the knowledge problems formulated in the research questions. Firstly, an introduction to the basic concepts surrounding routing problems is given, together with two examples that closely resemble the situation in Amsterdam (Section 3.1). Secondly, different solution approaches to routing problems are examined to give an overview of the possible approaches (Section 3.2). Subsequently, the Inventory Routing Problem (IRP) is discussed more extensively (Section 3.3) as this is the routing problem that most closely resembles the situation in Amsterdam. After that, the literature specifically researching the improvement of waste collection planning is studied (Section 3.4). To be able to evaluate different planning methodologies, we study different modeling approaches that can be used to model the waste collection system of Amsterdam (Section 3.5). Finally, the findings of the literature review are summarized in a conclusion and the contribution of this thesis to the existing literature is noted (Section 3.7).

3.1 - Introduction to routing problems

This section aims to give a general introduction to several concepts surrounding routing problems and their practical applications. We start by discussing an important input factor to all routing problems: shortest path routing. After that, a general classification of routing problems is discussed with practical examples. Subsequently, two specific routing problems are introduced: the Vehicle Routing Problem (VRP) and Inventory Routing Problem (IRP).

3.1.1 - Shortest path routing

One of the prerequisites for solving routing problems is being able to compute the shortest path between all locations that are considered in the routing problem in question. Shortest path routing is used to generate cost matrices, often based on distance, time, or a combination of the two, between locations (Ticha et al., 2017). One of the leading uses of shortest path routing is in transportation, but it also has applications in artificial intelligence, operations research, and computer science (Tommiska & Skytta, 2001) (Huang, Wu, & Zhan, 2007), making it a popular research topic.

A classic approach to solve the shortest path problem is Dijkstra’s algorithm (Dijkstra, 1959). However, since its publication in 1959, the performance of shortest path algorithms has improved drastically, leading to algorithms that are up to three million times faster than Dijkstra’s original algorithm (Delling et al., 2009). Other well-known algorithms include the A* search (Goldberg, Kaplan, & Werneck, 2006) and Bellman’s algorithm (Bellman, 1958).

Traditionally, most literature focuses on a single-objective shortest path. However, Tarapata (2007) argues that in many practical applications, multicriteria shortest path problems are more suitable. This is also the case in Amsterdam, where the municipality prescribes three criteria of importance: costs, emission, and nuisance.

3.1.2 - General types of routing problems

Routing problems can be classified as either node- or arc routing problems (Pearn, Assad, & Golden, 1987). This classification depends on where the demand is located in the underlying network. In node routing problems, the demand exists at the customers who are represented by nodes. In contrast, the demand in arc routing problems occurs on the arcs that can be traversed. The difference between arc routing and node routing is illustrated in Figure 7.

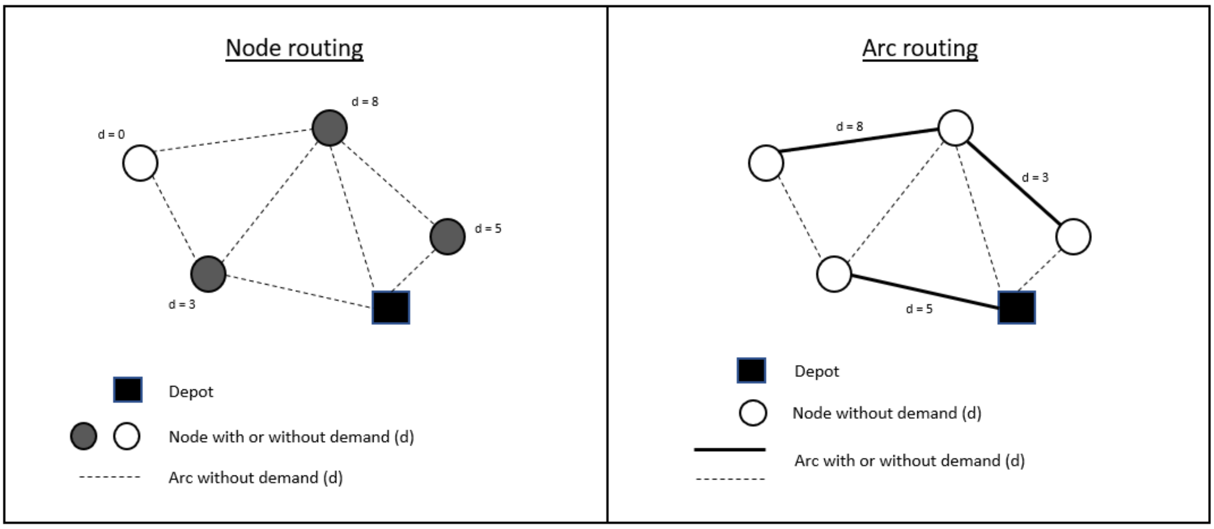

Figure 7: Node versus arc routing problems

A commonplace example of a node routing problem is the Travelling Salesman Problem (TSP). The general TSP can be formulated as follows: A salesman wants to find the shortest route in which to visit all customers and return back home (Lenstra & Kan, 1975). A generalization of the TSP is the VRP, which was first introduced by Dantzig and Ramser (1959). The VRP is a combinatorial optimization problem that is used to determine the optimal route to deliver a given set of orders to a given set of customers from one or multiple depots (Kumar & Panneerselvam, 2012). The VRP is discussed further in the next section.

Most research into arc routing problems is done on the specific case of the Capacitated Arc Routing Problem (CARP). This problem arises when streets have to be traversed, for example for maintenance, snow removal, or road gritting. Each road (i.e., arc) has to be serviced by exactly one vehicle and all vehicles have a limited capacity (Hertz, Laporte, & Mittaz, 2000).

Despite the difference between node and arc routing problems, they are of the same complexity. As for each arc routing problem, there exists an equivalent node routing problem and vice versa (Oppen & Løkketangen, 2006). Nevertheless, research into solving node routing problems is much more prevalent. Because of this, there exist several problems for which the computational results of the node routing version outperform their arc routing counterparts, for example, the Capacitated VRP (Baldacci & Maniezzo, 2004). Therefore, especially transformation techniques from arc to node routing problems have gotten some attention in literature (Pearn, Assad, & Golden, 1987) (Baldacci & Maniezzo, 2004).

Nevertheless, there are also instances where arc routing techniques are used in a node routing environment to reduce the problem size. Oppen and Løkketangen (2006) show that the problem size decreases and solution quality increases when customers are aggregated into clusters, represented by arcs. However, they also note that it is only applicable in specific cases where there are a high number of customers on so-called road segments, such as is the case with house-to-house waste collection or mail delivery.

The problem faced in Amsterdam can be classified as a node routing problem. The demands occur at the containers which do not adhere to the required characteristics of arcs in an arc routing problem. Therefore, the remainder of this section focuses on two specific node routing problems that most closely resemble the situation in Amsterdam: the VRP and IRP.

3.1.3 - Vehicle Routing Problem

The VRP was first introduced by Dantzig and Ramser (1959), who named it the "truck dispatching problem". The VRP is a generalization of the TSP, which adds the condition that dictates specific deliveries to be made to specific locations (Dantzig & Ramser, 1959). The most general formulation of the VRP concerns itself with constructing optimal routes between one depot and a number of customers that have a known demand (Laporte, 1992).

Since its inception in 1959, the VRP has become one of the most extensively studied combinatorial optimization problems (Cordeau et al., 2002). Due to this popularity, the original VRP has been extended by adding real-life characteristics to the problem (Braekers, Ramaekers, & van Nieuwenhuyse, 2016). However, most literature is limited to studying one additional characteristic at a time, disregarding that real-life cases should often adhere to a wide range of characteristics simultaneously (Braekers, Ramaekers, & van Nieuwenhuyse, 2016). Some variants with relevance to the problems faced in Amsterdam are listed and elaborated briefly below in Table 5.

Table 5: Variants of the Vehicle Routing Problem (adapted from Braekers, Ramaekers, & van Nieuwenhuyse (2016))

Characteristics	Description
Capacitated	In the capacitated VRP all vehicles have a limited capacity. This capacity is often expressed in weight, volume, or both simultaneously.
Stochastic demand	Instead of known, deterministic demand, the demand is unknown and stochastically distributed.
Satellite facilities	Satellite facilities allow vehicles to replenish their load during a route, without having to visit a depot location, enabling them to continue serving customers
Multiple depots	Instead of one depot from which customers can be serviced, a multi depot problem has multiple depots.
Heterogeneous vehicles	The general VRP has a homogeneous fleet which means that all vehicles are identical. The VRP with heterogeneous vehicles has different types of vehicles, each with its own characteristics.
Time-dependent travel times	Instead of fixed travel times between locations, the travel times are dependent on the time of the day. For example, to account for the increase in travel time during rush hours.
Dynamic	In dynamic VRPs the information is revealed continuously and partially after the schedule has already been constructed.

3.1.4 - Inventory Routing Problem

The IRP is an important and one of the most challenging extensions of the VRP (Bertazzi, Savelsbergh, & Speranza, 2008). Its pioneering paper was written by Bell et al. (1983) who focus on the efficient distribution of industrial gases to customers. This is an early example of a common practical situation in which the IRP arises, vendor managed inventory (VMI) (Coelho, Cordeau, & Laporte, 2014). VMI is a modern supply chain strategy in which the supplier, or vendor, becomes responsible for the inventory management of its customers (Sari, 2007). This means that the supplier is able to schedule deliveries to its customers itself, instead of being directed by the customer's order. This allows the supplier to potentially benefit from planning its deliveries in such a way that adjacent customers can be serviced together on the same day, saving travel costs. Accordingly, the objective of the IRP is to minimize the delivery costs, while ensuring no customer experiences stock-outs (Bard et al., 1998).

The IRP combines two theoretical fields: inventory management and routing. During the IRP, inventory management decisions are taken such as: when should we serve each customer and how much should we deliver each customer, simultaneous with routing decisions such as: how should the chosen customers be combined in vehicle routes (Coelho, Cordeau, & Laporte, 2014). This connotes that the decisions should be taken in two dimensions: the time and space dimensions. This creates an extra level of complexity in comparison to traditional routing problems, where the decision maker is only concerned with the space dimension.

The integration of these decisions is important, as the results of one field influence the possibilities in the other. For example, timing decisions such as the inventory allocation dictates which customers are served on the same day, restricting the routing decision which is related to the space dimension of the IRP. The objective of the IRP is to minimize the delivery costs while ensuring no customer experiences stock-outs (Bard et al., 1998). A more general formulation of the IRP is given by Campbell et al. (1998): "The IRP is concerned with the repeated distribution of a single product from a single facility to a set of N customers over a given planning horizon of length T, possibly infinity. […] The objective is to minimize the average distribution costs during the planning period without causing stockouts at any of the customers".

In addition to the earlier mentioned fields of inventory management and routing, Baita et al. (1998) introduce another aspect of the IRP: dynamicity. This refers to the dynamic framework in which the inventory and routing decisions are taken; decisions are taken at different times, where earlier decisions significantly impact later decisions (Baita et al., 1998). This is repeated by Moin and Salhi (2007), who describe the IRP as a medium-term problem, in contrast to the short-term character of the VRP. For example, the decision to postpone serving a customer may seem beneficial in the short term, but this customer cannot be postponed indefinitely. One of the most important challenges posed by the IRP is finding the optimal timing of deliveries considering both the time and space dimensions.

From this, the similarity between the IRP and the waste collection planning of Amsterdam is clear. In Amsterdam, the municipality is responsible for determining the emptying time of each container, which are the customers in this analogy with the IRP. Moreover, the municipality objective is to minimize collection costs, while ensuring a high service level, identical to the objectives of a general IRP. Because containers are visited not to distribute something, as in the general IRP, but to collect something, the problem is known as the reverse IRP (Mes, Schutten, & Rivera, 2014).

The IRP is an NP-hard problem because it can be reduced to the classical VRP (Coelho, Cordeau, & Laporte, 2014). This means that there is no known algorithm that can solve the IRP in polynomial time, but only in nondeterministic polynomial time (Papadimitriou, 2003). This makes solving the problem to optimality computationally infeasible as the problem size increases (Woeginger, 2003).

As mentioned earlier, Section 3.3 elaborates further on the IRP and its variants, but first, we introduce several general solution approaches for routing problems in the following section.

3.2 - Introduction to solution techniques

Finding solutions for these types of routing problems has been the focus of an abundance of literature. This section aims to give a short overview of the possible techniques that are used. We use the classification of Coelho, Cordeau, and Laporte (2014) who classify solution techniques for the IRP into three categories: exact methods, heuristics, and metaheuristics.

3.2.1 - Exact methods

The aim of exact methods or algorithms is to solve the problem to proven optimality. However, as the IRP is NP-hard this becomes infeasible even for fairly smaller instances. Nevertheless, exact techniques are occasionally used to solve at least parts of the IRP. Two techniques that are used more frequently are described by Woeginger (2003): dynamic programming and branch-and-bound algorithms.

Dynamic programming involves dividing large complex problems into smaller subproblems that are more easily solvable. By using the solutions to these subproblems, dynamic programming is able to solve the larger, complete problem (Woeginger, 2003).

A more commonly used exact approach in the IRP literature are branch-and-bound or branch-and-cut algorithms. Both approaches are based on search trees, where the solution space is partitioned into smaller subsets for which all feasible possibilities, represented by 'branches', are evaluated (Woeginger, 2003). The branches for which it can be proven that they cannot result in an optimal solution, even before the complete solution space is filled, can be 'pruned'. This branch is then discarded and not explored further, reducing the computation time.

3.2.2 - Heuristics

Because solving a problem to optimality using exact methods is not always feasible, as is the case with the IRP, a lot of research focuses on approaching the optimum using approximation heuristics that result in acceptable solutions with less computation time (Nilsson, 2003). Blum and Roli (2003) distinguish two basic categories of heuristics: constructive methods and local search methods.

Constructive methods are designed to create a feasible solution from scratch. They are typically fast, but also provide relatively poor solutions (Blum & Roli, 2003). Examples of construction heuristics used for routing problems are the: Nearest Neighbor, Greedy, and Christofides heuristics (Nilsson, 2003).

Local search methods start with a feasible solution and attempt to iteratively improve that solution (Blum & Roli, 2003). These methods are called local search because the improvements are sought in 'neighboring' solutions in the solution space. An example of such a heuristic for routing problems is the k-opt algorithm (Nilsson, 2003).

3.2.3 - Metaheuristics

Recently, an increasing amount of research is done into metaheuristics (Coelho, Cordeau, & Laporte, 2014). Metaheuristics are high-level heuristics that 'guide' the search process to efficiently search the solution space for a near-optimal solution (Blum & Roli, 2003). An advantage of metaheuristics is that they do not get stuck in local optima, as can be the case with generic heuristics. Examples of these kinds of metaheuristics are: Greedy Randomized Adaptive Search Procedure (GRASP), Simulated Annealing, Tabu search, and Ant Colony Optimization (Blum & Roli, 2003).

3.3 - Inventory Routing Problems

To systematically review the current literature on the IRP, a classification matrix is drafted. Using a classification matrix ensures that the literature review is performed in a concept-centric manner, i.e., structured around concepts instead of individual authors (Webster & Watson, 2002).

Several classification schemes from existing literature reviews covering the IRP are studied to identify relevant concepts that can be used to characterize IRP papers. Moin and Salhi (2007) categorize papers according to the modeled time horizon: single period, multiperiod, and infinite models. Stochastic models are discussed separately as a relatively newer area of research. However, most reviews classify IRP papers based on their practical characteristics. Examples are the reviews of Baita et al. (1998), Andersson et al. (2010), Bertazzi and Speranza (2012), and Coelho, Cordeau, and Laporte (2014). A complete overview of their classification schemes can be found in Appendix 4. Based on the classification elements used in these papers in combination with relevant aspects relevant from the context of Amsterdam we formulate the following classification structure, see Table 6.

Table 6: Classification structure based on Baita et al. (1998), Andersson et al. (2010), Bertazzi and Speranza (2012), and Coelho, Cordeau, and Laporte (2014)

Classification element	Attribute	Alternatives		
Structure	Topology	One-to-many	Many-to-many	Others
	Number of items	Single	Multiple	
Time horizon	-	One period	Multiperiod	Infinite
Demand	-	Deterministic	Stochastic	
Fleet	Composition	Homogeneous	Heterogeneous	
	Size	One	Multiple	Infinite
Solution approach	-	Various		

The first classification element, the problem structure, has two aspects: topology and number of items. The topology refers to the relationship between the number of depots and customers. We distinguish one-to-one, one-to-many, and many-to-many relationships. The second aspect of the problem structure is the number of items, e.g., different products or waste fractions, considered in the problem. Problems consider either one single product or multiple products. The next element is the time horizon that is considered in the problem. We categorize three different time horizons: one period, multiperiod, or infinite horizon models. Papers can also be classified according to the displayed demand characteristics, multiple characteristics are suggested, such as seasonality, et cetera. However, we choose to only consider the uncertainty of demand: deterministic or stochastic. Next, we differentiate papers according to the characteristics of the used fleet of vehicles. Vehicle fleets can be composed of one vehicle, multiple but limited, or an infinite number of vehicles. In the case of multiple vehicles, the fleet can be homogeneous or heterogeneous. In a homogeneous fleet, all vehicles are identical, while there are differences between vehicles in a heterogeneous fleet. The last classification element is the solution approach. This describes the approach that is used to solve the IRP. The last column of the concept matrix is reserved for further comments on aspects that do not immediately fit in the structure of the concept matrix, but are of note. The complete concept matrix is shown in Table 7.

Table 7: Classification matrix of the Inventory Routing Problem

Paper[1]	Problem structure		Time horizon	Demand	Fleet		Solution approach[2]	Additional notes
	Topology	No. of items			Composition	Size		
[1]	One-to-many	Single	Single period	Stochastic	Homo-geneous	Multiple	Decomposition and MIP	Treat IRP as variant of VRP
[2]	One-to-many	Single	Multi-period	Deterministic	Homo-geneous	Multiple	Various	Comparison of computational algorithms
[3]	One-to-many	Single	Multi-period	Stochastic	Homo-geneous	Multiple	Decomposition: (i) clustering, (ii) routing	Rolling horizon, satellite facilities
[4]	One-to-many	Multiple	Multi-period	Stochastic	-	One	Bi-level iteration heuristic with IP and TSP	Proof of lower bound calculations
[5]	One-to-many	Single	Multi-period	Stochastic	Homo-geneous	Multiple	Dynamic programming approximation methods	IRP as Markov decision process, direct deliveries
[6]	One-to-many	Single	Multi-period	Deterministic	Homo-geneous	Multiple	Decomposition to (i) IP for high level plan and (ii) heuristics for detailed plan	Rolling horizon, decomposition in the time dimension
[7]	One-to-many	Single	Multi-period	Deterministic	Homo-geneous	Multiple	Decomposition to (i) Inventory Allocation, (ii) TSP, and (iii) VRP	Variable lead times due to traffic congestion
[8]	?-to-many	Single	Multi-period	Deterministic	Homo-geneous	Infinite	Hybrid heuristic combining Tabu search and SA	Solution method tested using simulation, adds variable depot locations
[9]	One-to-many	Single	Multi-period	Deterministic	-	One	MILP relaxation, branch-and-cut algorithm	First exact algorithm, based on valid inequalities
[10]	Many-to-many	Single	Multi-period	Deterministic	Homo-geneous	Multiple	Integer Programming and local search heuristic	Specific attention for problem size reduction
[11]	One-to-many	Single	Multi-period	Deterministic	-	One	ALNS algorithm and network flow algorithm	IRP with transshipment
[12]	Many-to-many	Multiple	Multi-period	Deterministic	Homo-geneous	Multiple	MILP model	Tested on both 4 randomly generated as 1 real-life case
[13]	One-to-many	Multiple	Multi-period	Deterministic	Homo-geneous	Multiple	Exact solution with branch-and-cut algorithm	Feasible up to medium size instances of several problem classes
[14]	One-to-many	Single	Infinite (cyclic)	Deterministic	Homo-geneous	Infinite	ALNS algorithm	Periodic IRP
[15]	One-to-many	Multiple	Multi-period	Deterministic	Homo-geneous	Multiple	Three phase heuristic: (i) replenishment plan, (ii) sequencing, (iii) planning and routing with MILP	Includes Lagrangian-based heuristic in first phase

Table 7: Concept matrix of the Inventory Routing Problem (continued)

Paper[1]	Problem structure		Time horizon	Demand	Fleet		Solution approach[2]	Additional notes
	Topology	No. of items			Composition	Size		
[16]	One-to-many	Multiple	Multi-period	Deterministic	Homo-geneous	Infinite	Hybrid GA based on decomposition to (i) allocation and (ii) routing	Accounting for carbon emission regulations
[17]	One-to-many	Single	Infinite (cyclic)	Deterministic	Homo-geneous	Multiple	Decomposition to (i) routing and (ii) scheduling	Proposes Vehicle Decrease Heuristic
[18]	One-to-many	Single	Multi-period	Deterministic	Hetero-geneous	Multiple	Branch-and-cut algorithm	Introduces green-IRP. Analyses impact of heterogeneous vehicles
[19]	One-to-many	Multiple	Multi-period	Deterministic	Homo-geneous	Multiple	Hybrid RVND	Case study with planned transshipment
This thesis	One-to-many	Single	Multi-period	Stochastic	Homo-geneous	Infinite	Decomposition to (i) selection, (ii) day assignment, (iii) routing, see Chapter 4	Satellite facilities, integrated decision on time and space dimension

1 With: [1] Federgruen & Zipkin (1984) , [2] Dror, Ball, & Golden (1985), [3] Bard, Huang, Jaillet, & Dror (1998), [4] Qu, Bookbinder, & Iyogun (1999), [5] Kleywegt, Nori, & Savelsberg (2002), [6] Campbell & Savelsbergh (2004), [7] Chiou (2005), [8] Liu & Lin (2005), [9] Archetti, Bertazzi, Laporte, & Speranza (2007), [10] Savelsbergh & Song (2008), [11] Coelho, Cordeau, & Laporte (2012), [12] Ramkumar, Subramanian, Narendran, & Ganesh (2012), [13] Coelho & Laporte (2013), [14] Aksen, Kaya, Salman, & Tüncel (2014), [15] Cordeau, Laganà, Musmanno, & Vocaturo (2015), [16] Cheng, Qi, Wang, & Zhang (2016), [17] Chitsaz, Divsalar, & Vansteenwegen (2016), [18] Cheng, Yang, Qi, & Rousseau (2017), and [19] Peres, Repolho, Martinelli, & Monteiro (2017).

2 MIP = Mixed Integer Programming, IP = Inventory Problem, SA = Simulated Annealing, MILP = Mixed Integer Linear Programming, ALNS = Adaptive Large Neighborhood Search, RVND = Randomized Variable Neighborhood Descent.

3.3.1 - Problem structure

The majority of reviewed IRP literature considers networks with a one-to-many topology supplying a single product. A one-to-many topology occurs in the case where a single, often central, facility or depot services a set of customers (Andersson et al., 2010). Savelsbergh and Song (2008) note that, even when multiple depots are considered, customers are almost always assigned to a single depot, decomposing the problem into multiple one-to-many problems. It is argued that, because of real-life complexities, such as insufficient production capacity at the supplier, this approach is not always feasible to solve real-life problems. To overcome the limitations presented by the regular one-to-many IRP, Savelsbergh and Song (2008) propose the IRP with continuous moves, IRP-CM, that allows for multi-day routes. A different type of topology is presented by the Combined Location Routing and Inventory Problem (CLRIP). In this problem class, the location of the depots are not given, but present additional decision variables. Liu & Lin (2005) propose a decomposition heuristic for the CLRIP that first solves the depot location-allocation problem and subsequently solves the IRP.

An extension to the traditional topology, with depots and customers, is to add locations with other functionalities such as satellite facilities or by allowing transshipments. Bard et al. (1998) consider the IRP with satellite facilities (IRPSF). Satellite facilities function as additional depots, allowing vehicles to replenish their inventory during a route. Strategically located satellite facilities prevent necessary trips back to the central depot to restock, in turn effectively increasing vehicle capacity (Bard et al., 1988). Bard et al. (1998) propose a decomposition approach to solve the IRPSF that starts by selecting and assigning customers to days and then solving the resulting VRP with satellite facilities (VRPSF) for all days. Several heuristics for the VRPSF are considered: a revised Clark & Wright (C&W) algorithm, a revised sweep algorithm, and a GRASP heuristic. The effectiveness of these solutions is tested on randomly generated problem instances, showing the C&W algorithm outperforming the other algorithms slightly (Bard et al., 1998). Coelho, Cordeau, and Laporte (2012) introduce the IRP with transshipment (IRPT), where products can be transported either from supplier to customer or from customer to customer. An ALNS heuristic, in combination with a network flow algorithm, is proposed to solve the IRPT (Coelho, Cordeau, & Laporte, 2012).

While most literature considers the distribution of a single product, Coelho and Laporte (2012) and Cordeau et al. (2015) extend this to a multi-product, multi-vehicle IRP (MMIRP). Although both papers attempt to solve the same problem, widely different solution approaches are used. Coelho and Laporte (2013) devised a branch-and-cut algorithm to give the exact solution. Whereas Cordeau et al. (2015) take a three-step decomposition approach: (1) constructing delivery plans, (2) determining delivery sequences, and (3) a re-optimization phase that used a MILP model to improve the solution. Ramkumar et al. (2012) combine a many-to-many topology with multiple different products and propose a MILP for the problem. Their solution is tested on four randomly generated datasets and one real-life case study which resulted in a cost reduction of more than 7% to the total current costs.

3.3.2 - Time horizon

The modeled time horizon dictates over which period of time the problem is optimized. Single period models solely focus on minimizing the costs over one period, while multi-period models consider the costs over a longer horizon (Moin & Salhi, 2007). While single period models are less complex than their multi-period counterparts, they often offer worse solutions. This is caused by the short-term approaches' tendency to postpone as many deliveries as possible, which is beneficial in the short term, but has negative effects on the long-term planning (Campbell et al., 1998) (Moin & Salhi, 2007). Because of this, contemporary literature has a clear inclination towards multi-period models.

However, especially earlier contributions towards IRP literature still consider the single-period IRP, such as Federgruen and Zipkin (1984), one of the first papers to combine routing and inventory decisions. Federgruen and Zipkin (1984) attempt to minimize the combined transportation, holding, and shortage cost for one single period. Because just one period is considered, this problem can be seen as an extension of the VRP, which causes Federgruen and Zipkin (1984) to utilize many techniques also used in the VRP. The results of the combined approach were however much better than those acquired using regular VRP techniques: 6-7% cost reduction and 20% reduction of vehicles required (Federgruen & Zipkin, 1984).

To prevent the model from postponing as many deliveries as possible, the long-term effects of short-term decisions should be taken into account. One of the first approaches to do this was proposed by Dror, Ball, and Golden (1987), who introduced penalties and incentives to single-period models. Their approach models the multi-period IPR as consecutive single-period IRPs, in which additional cost factors are added to represent the expected future costs of decisions. These future costs account for, among others, the risk of postponing deliveries to almost stocked-out customers. Another way to model a multi-period IRP is to use a rolling horizon framework. For instance, Campbell and Savelsbergh (2004) implement a two-phase decomposition approach. During the first phase, a basic schedule is made for a longer k-day horizon, after which the second phase makes a complete planning for a shorter j-day horizon based on the results of the first phase. In this approach, the long-term costs are accounted for in the first phase, the long-term deliver schedule, and the short-term costs are minimized during the second phase, which mainly involves routing.

Another way to account for longer-term costs is to by using an infinite planning horizon. Models with infinite planning horizon often aim to minimize the long-run or mean average of all costs (Moin & Salhi, 2007). Moin and Salhi (2007) also note that infinite horizon models are often based on fixed-partition policies. Fixed partition policies separate all customers into different customer sets. These sets are then served independently from each other at their collective optimal replenishment rate (Anily & Bramel, 2004). One of the problems with fixed-partitioning is that it does not easily allow for coordination between clusters where this may be beneficial. An example of fixed-partitioning is given by Chitsaz, Divsalar, and Vansteenwegen (2016), who assign customers to trips and then determine the optimal cycle time for each trip. This presents another variant of the IRP, the Cyclic Inventory Routing Problem (CIRP). Each partition is served with a constant replenishment interval, making the time horizon infinitely long. A slightly different approach is introduced by Aksen et al. (2014), who solve the IRP as a Periodic Inventory Routing Problem (PIRP), in which a predetermined planning repeats itself each time period, for example, weekly.

3.3.3 - Demand characteristics

An important modeling factor of the IRP are the demand characteristics of the customers. Because of the complexity inherent to the IRP, most papers only consider a deterministic demand pattern, where all demand is known beforehand. However, Moin & Salhi (2007) remark that, in reality, customer demand oftentimes has a stochastic nature. To capture this stochasticity, the stochastic IRP models customer demand using a predetermined probability distribution.

Kleywegt, Nori, and Savelsbergh (2002) model the stochastic IRP as a Markov decision process. The current inventories at customers are defined in a state space. Transitions to other states are given by a Markov transition function which is governed by a known joint probability distribution of customer demands (Kleywegt, Nori, & Savelsbergh, 2002). As the number of states grows exponentially with the number of customers, Kleywegt, Nori, and Savelsbergh (2002) propose approximation methods to solve the Markov decision process. Qu, Bookbinder, and Iyogun (1999) choose to model customer demand in the form of a Brownian motion process. However, because they assume vehicle capacity to be infinite, this only affects the inventory management side of their problem decomposition.

Research has also given attention to different distribution policies that can be used to replenish inventories. The most common are the order-up-to and the maximum level policies (Bertazzi & Speranza, 2012) (Coelho, Cordeau, & Laporte, 2012). A modified version of another alternative, the periodic review policy, is used by Qu, Bookbinder, and Iyogun (1999) to complement the cyclic nature of their IRP solution approach.

3.3.4 - Vehicle fleet

The available vehicle fleet dictates the number and characteristics of the vehicles that can be used to service customers in the IRP. We distinguish two important vehicle fleet characteristics: size and composition. Both characteristics affect the complexity of the IRP and the possible solution approaches. For example, Archetti et al. (2007) simplified the IRP to a single vehicle fleet, which enabled them to develop an exact branch-and-cut algorithm. However, they do note that in reality this simplification would often be infeasible. In cases where the number of vehicles is unrestricted the vehicle fleet is called infinite. Cheng et al. (2016) use an infinite fleet size to increase the flexibility to choose the required number of vehicles later. However, in most practical scenarios, a limited number of vehicles is often available. Therefore, most IRPs consider the case where multiple, but a limited amount of vehicles are available.

An important distinction in multi-vehicle fleets is if they are homogeneous or heterogeneous. Homogeneous vehicle fleets consist of identical vehicle, while heterogeneous fleets have vehicles that differ in, for example, capacity and functionalities. Almost all papers consider homogeneous vehicle fleets, but Cheng et al. (2017) consider a fleet composed of light-, medium-, and heavy-duty vehicles. Here, the heterogeneity stems from the different capacities of the vehicles. Using a branch-and-cut algorithm, Cheng et al. (2017) show that using such a heterogeneous fleet offer significant reductions in costs in comparison to homogeneous vehicle fleets of one of the three vehicle types.

3.3.5 - Solution approaches

Coelho, Cordeau, and Laporte (2014) identify two general solution approaches to the IRP, as also noted in Section 3.2: exact approach and (meta)heuristic approaches. However, Andersson et al. (2010) and Coelho, Cordeau, and Laporte (2014) both note that, because of the problem complexity, only small instances can be solved to optimality. Therefore, most solution approaches proposed in literature use heuristics, metaheuristics, or mathematical programming techniques which are ended before proven optimality (Andersson et al., 2010).

The majority of exact methods proposed to solve the IRP are based on branch-and-cut approaches. Archetti et al. (2007) were the first to propose such an approach for the single-vehicle IRP. They proposed a mixed-integer programming model and derived a new set of valid inequalities strengthening the formulation of the linear relaxation. With this approach, Archetti et al. (2007) solve instances with up to 50 customers to optimality with a small time horizon of three periods. An extension on Archetti et al. (2007) is proposed by Coelho and Laporte (2013), who extend the branch-and-cut algorithm to include the multi-product multi-vehicle IRP. While this increases the complexity of the problem, their proposed solution is able to solve larger problem instances to optimality than Archetti et al. (2007). However, this is at the expense of the required computation time which often exceeds 1 hour. An approach to limit the required computation time is to stipulate a maximum run time for the exact solution method. Although, even in papers using such methods, such as Cheng et al. (2017) and Ramkumar et al. (2012), the stipulated time limits are very high, respectively 1 and 8 hours in these papers. Cheng et al. (2017) do not that their algorithm often finds satisfactory solutions within 5 minutes, but this is limited to problem instances of 20 customers.

Due to the often large problem instances encountered in real-life, most proposed solution approaches are heuristics (Coelho, Cordeau, & Laporte, 2014). A high number of heuristic approaches are based on a certain decomposition approach which decomposes the problem into several smaller subproblems for which less complicated solution approaches can be applied. The decomposition method, however, varies from paper to paper. Moin and Salhi (2007) observe that most papers adopt an approach that decomposes the IRP into two separate problems: the inventory and travelling salesman problem. This approach can either be designed as (1) inventory-first, route-second, where the inventory management problem is solved first and the routes are based on the resulting customer clusters or (2) route-first, inventory-second, where routes are found first after which the IRP formulation is completed by, for example, using linear programming techniques (Moin & Salhi, 2007). Baita et al. (1998) identify a third recurring decomposition approach: decomposing over the time horizon. In this approach, a long-term IRP can, for example, be decomposed into multiple single period problems. More specific examples of decomposition approaches are the three step decomposition by Bard et al. (1998): customer selection, day assignment, and routing; Cordeau et al. (2015): constructing replenishment plans, sequencing, planning and routing; and the two phase algorithm of Chitsaz, Divsalar, and Vansteenwegen: routing and scheduling.

3.4 - Waste collection planning

Because of the growing societal attention to sustainability and the realization that advances can be made in the efficiency of waste collection and management, it is a growing topic in literature (Beliën, de Boeck, & van Ackere, 2014). In this short literature review, we discuss the trend of current literature, several approaches to the problem of waste collection planning, and new technologies impacting the possibilities of waste collectors.

Currently, a large portion of the literature focusses on the collection of residential, house-to-house collection that is best represented by arc routing problems (Beliën, de Boeck, & van Ackere, 2014). Moreover, the papers that do discuss the collection of larger containers, most approach the problem exclusively as a VRP, or as often called a Waste Collection Vehicle Routing Problem (WCVRP) (Beliën, de Boeck, & van Ackere, 2014) (Han & Ponce-Cueto, 2015). There are however also other approaches of which we discuss several below.

The most common approach is to model the waste collection as a WCVRP. WCVRPs are concerned with finding the optimal route to collect waste from a set of containers. Collection vehicles leave the depot empty, collect waste from containers and unload their collected waste at a disposal facility the end of the route, or when necessary. At the end of the day, the collection vehicle returns to the depot (Benjamin & Beasley, 2010). This formulation implies that the set of containers that is to be collected is already known. To solve the WCVRP, a broad range of solution approaches is applied (Beliën, de Boeck, & van Ackere, 2014). Two examples are the use of Solomon's insertion algorithm by Kim, Kim, and Sahoo (2006) and the application of different metaheuristics, such as Tabu search and variable neighborhood search by Benjamin and Beasley (2010), both of which are also commonly used in VRP literature.

Two less common approaches are to model the waste collection process as a Team Orienteering Problem (TOP) or an IRP. Both these variants, unlike the WCVRP, include the decision of which containers should be emptied, albeit in widely different ways. In the TOP, each customer has an associated profit and routes are constructed to maximize the total profit, while complying with restrictions such as route length and duration (Ferreira et al., 2014). The profit associated with waste containers can then, for example, be a representation of its priority. Ferreira et al. (2014) developed a genetic algorithm to solve the TOP which was competitive on most benchmark instances. The IRP is already extensively discussed earlier in this chapter, as it is also the approach taken in this thesis. Mes et al. (2014) model the waste collection problem as a reverse-IRP and use a combination of optimal learning and simulation to tune their heuristic, which selects containers based on a must- and may-go classification.

The increased attention to the improvement of waste collection and management is not limited to literature discussing routing problems. It is also discussed as part of another upcoming trend: smart cities (Medvedev et al., 2015). Medvedev et al. (2015) discuss the potential of using new advancements such as the Internet of Things to enable dynamic waste collection, supporting both the decision when containers should be emptied and what routes truck should follow.

3.5 - Modeling approaches

Before implementing a new planning methodology, the expected effects of the new methodology on the planning performance should be evaluated. This can be done by monitoring the system, the waste collection of Amsterdam, under different circumstances and comparing the differences. Law (2015) identifies three modeling approaches to study systems: physical models, analytical models, and simulation models.

Testing with physical models involves altering the real, physical system, or a scale model of the physical system, and evaluating the effects of the interventions directly (Law, 2015). A benefit of this modeling approach is that, because the real system is altered, the model's validity is certain. However, altering the physical system is often infeasible as this is time-consuming, expensive, or dangerous.

Analytical and simulation models are both examples of mathematical models. With mathematical models, it is not required to modify the actual system risking productivity loss or dangerous situations. Instead, a mathematical model represents the system in terms of quantitative and logical relationships between the system's entities (Law, 2015). To test the effects of changes to the system, these relationships can then be altered to represent the interventions and the performance of the system before and after can be compared. Analytical models are better suitable to study more elementary systems, while simulation models are better equipped to deal with larger, complex systems (Law, 2015). Because of the complexity of the waste collection environment and the added stochasticity of unpredictable waste deposit patterns, we therefore choose to model the waste collection system using a simulation modeling approach. Simulation is a common method to model the complexities of waste collection systems and is often used in literature (Mes, Schutten, & Rivera, 2014) (Johansson, 2006) (Das & Bhattacharyya, 2015).

3.6 - Conclusion

This chapter gives a broad overview of the literature related to the waste collection problem faced by Amsterdam. The aim of this chapter is to answer research question 2: What relevant routing problems and corresponding solution approaches are described in literature? To do this, literature is discussed to introduce the relevant routing problems and solution techniques. Moreover, an extended review is done on the literature directly related to the IRP, and approaches taken to specifically improve waste collection planning are discussed. This chapter is concluded by discussing the applicability of the literature to the problem of Amsterdam and noting the contribution of this thesis to the existing literature.

3.6.1 - Applicability to research

During the literature review, several different types of routing problems are identified and discussed. It is recognized that node routing, and more specifically the IRP, best resemble the waste collection planning problem faced by Amsterdam. This conclusion is drawn based on the geographical characteristics of containers, which indicate a node routing problem, and the identical decisions and objectives that are associated with both the IRP and the waste collection planning of Amsterdam. Mes et al. (2014) differentiate the application of IRP to waste collection planning from the general IRP. They argue that, in waste collection, as waste collected instead of distributed, we are facing the reverse IRP. However, they also note that solutions applied to the IRP are also applicable to the reverse IRP (Mes et al., 2014). Therefore, this thesis can still learn from approaches applied to the general IRP.

Because of the wide range of applications and simplifications of real-life problems, there exists a great number of variants to the general IRP, as discussed in Section 3.3. The most important differentiation between the majority of literature and the problem faced by Amsterdam is the topology of the waste collection system. Where most IRPs are limited to a one-to-many topology, where one depot is used to serve multiple customers, the topology of Amsterdam is more complex. This is caused by the inclusion of four different locations types, all with different functions: containers, wharfs, satellite facilities, and waste processing facilities. However, even though the problems discussed are not identical, lessons can still be learned from approaches to similar problems. For example, about the applied decomposition approaches and heuristics.

The applicability of the literature discussing improvement in the waste collection planning is limited. Most papers apply the WCVRP, which neglects an important decision that has to be taken during the waste collection planning: the timing of emptying a container (Beliën, de Boeck, & van Ackere, 2014). Moreover, for our purposes, modeling the waste collection planning as a Team Orienteering Problem (TOP), as proposed by Ferreira et al. (2014), also seems unsuitable. This is the case because the objective of the TOP is to maximize profit every single day, which would imply that vehicles should always be used to maximal capacity. However, this would likely result in emptying containers more frequently than necessary.

Next to research on literature directly related to waste collection planning, we also researched the different modeling approaches that are used to evaluate and compare different planning methodologies. Three options are identified: physical, analytical, and simulation models (Law, 2015), of which the latter seems best suited for the context of this thesis.

3.6.2 - Contribution to existing literature

The contribution of this thesis to the existing literature is twofold: (1) the adjusted situation and specific conditions in which the IRP is applied and (2) the novel approach to solve the IRP.

The application of the IRP to the planning of waste collection is, as discussed earlier, limited. Most papers skip important decision steps in the planning of waste collection and others use different objectives. On the other side, literature discussing IRP is plentiful but lacks several critical characteristics that apply in Amsterdam, such as the complex topology.

Moreover, the majority of IRP literature utilizes some kind of decomposition approach. While there is nothing inherently wrong with that, the decomposition typically disconnects the time and space dimensions. This shortcoming prevents the decision maker to fully benefit from the characteristics of the IRP, which is created by the interconnectedness of the timing and routing decision.

The contribution of this thesis is to propose a novel solution approach that simultaneously considers the time and space dimension. Moreover, the solution approach is applied to a new situation, which is typically modeled using different theoretical problems, such as the WCVRP. The novel solution approach is elaborated in the following chapter, where it is also distinguished from other similar approaches.

Chapter 4 - Solution approach

This aim of this chapter is to give an outline of the proposed solution approach. However, firstly, a short recapitulation of the problem is given to show how the IRP of Amsterdam is related to other IRPs and introduce the relevant notations (Section 4.1). After that, several terms that are introduced in the proposed solution approach are explained beforehand (Section 4.2). Subsequently, the main steps of the solution approach are described (Section 4.3). Finally, the chapter is ended with a summarizing conclusion about the discussed solution approach (Section 4.4).

4.1 - Problem description

The objective of the waste collection planning is to empty all waste containers in time (i.e., before they overflow), with as little resources (e.g., vehicles, fuel, personnel) as possible. To do this, two decisions should be made: (1) when to empty each container and (2) how are the chosen containers scheduled in vehicle routes. In contrast with the traditional IRP that delivers goods to replenish customers, waste collection involves the collection of waste from containers. This variation of the IRP is called the reverse IRP.

In this problem, we have a set of containers C and a set of locations L. L is partitioned into four subsets representing a list of container locations $L^{cont} = \{l_1, \dots, l_n\}$ and three sets for additional locations L^{wh}, L^{wp}, and L^{sf}, respectively for the wharf, waste processor, and satellite facility locations. Naturally, every container $i \in \mathrm{C}$ corresponds with one container location in L^{cont}. The cost of travelling between two locations, i and j, are given for each pair of locations of L and is denoted as $c_{i,j}$. These costs are based on the Euclidean distances between two locations times a correction factor of 1,2 prescribed for urban areas by Levinson and El-Geneidy (2009).

Every day, the inhabitants of Amsterdam make deposits in the containers. As we work with discrete time intervals of one day, we assume that all waste deposits of the day are done at the end of the day, after the waste containers are emptied. The deposited amount per day is stochastic and we decided to model it using a Gamma distribution, given by $d_i \sim \mathrm{Gamma}\ (\mathrm{k}_i, \theta_i)$, where k_i and θ_i respectively denote shape and scale parameters. The Gamma-distribution is chosen to model the waste deposits as it is commonly used for this purpose (Mes, Schutten, & Rivera, 2014) and the available data is insufficient to reliably test the fit of other statistical distributions. The daily deposits result in the build-up of waste in each container at time t. Because the deposited amount of waste is stochastic, we differentiate between the expect amount of waste $\hat{u}_{i,t}$ and the actual amount of waste $u_{i,t}$ in container i at time t. All containers have a finite capacity to store waste w_i, which denotes the number of kilograms of waste that fits in container i. If a container overflows (i.e., if $u_{i,t} > w_i$), the municipality is notified and is obligated to empty the container the next day. Waste continues to be deposited at the overflowed container at the same rate as before. The amount of waste that exceeds the container's capacity is denoted by the deposit overflow at time t, $o_{i,t}$.

Containers should be scheduled to be emptied and routes are constructed determining in which order the containers are emptied. A route $r \in R$ is a list of containers that are visited sequentially. Each route starts at a wharf and ends with the combination of either a waste processor or satellite facility and finally a wharf. All waste is collected using a homogeneous vehicle fleet V, which is a simplification upon the actual situation in Amsterdam. Each route should have a designated vehicle to execute the route. Moreover, each vehicle has a finite capacity denoted by v_i. Vehicles can dispose of their waste load by visiting a waste processor or satellite facility, where the latter has an additional shipping penalty. The vehicle is then free to continue its route and empty more additional containers or return to a wharf. The scheduling decisions, i.e., which routes should be executed by which vehicles, are taken for each day t of a finite planning horizon of length T.

4.2 - Definitions in solution approach

This section introduces several new, less self-evident definitions that are used in the proposed solution approach and offers a short explanation.

Acceptable Overflow Probability	AOP	A general threshold value that specifies the acceptable probability that containers overflow.
Desired Emptying Day	DED_i	The expected day before which the DEL_i is reached.
Timing Penalty costs	$TP_{i,t}$	The additional costs of emptying container i either too early or too late on day t.
Cluster Fitness Approximation	$CFA_{i,j}$	The expected costs of adding container i to route or cluster j.
Penalty Emptying too Late	$PEtL_{i,t}$	The penalty for emptying container i too late, at time t.
Penalty Emptying too Early	$PEtE_{i,t}$	The penalty for emptying container i too early, at time t.
Overflow Probability	$OP_{i,t}$	The probability that the capacity of container i is exceeded at time t.
Single Container Route Costs	$SCRC_i$	A penalty measure that calculates the costs of making a route solely consisting of emptying container i, dumping the waste at the closest dump and returning to the wharf.
Expected Interval Length	EIL_i	The expected interval length between two visits of the container i.

4.3 - Proposed solution approach

This section presents the proposed solution approach which consists of three phases, as shown in Figure 8: container selection, day assignment, and route construction. The three phases are executed in a rolling planning horizon framework. This means that all phases are executed each day, while also considering the following days of the planning horizon in the decision-making process. An example is given in Figure 9, where a planning is made for a planning horizon of three days, but only the first day of the planning is executed. The entire planning is reconsidered again on the following day. Utilizing a rolling planning horizon enables us to evaluate the long-term consequences of decisions beyond the short-term impact on the current day. In doing so, the proposed solution approach is able to make decisions based on both time and space dimensions simultaneously. As only the first day of the planning is actually executed, Phase III only constructs routes for that particular day to save computation time.

The input for the proposed solution approach is twofold: (1) the container characteristics, which include the container locations, fill speeds, and expected fill levels, and (2) the experimental factors. These experimental factors are parameters that are used to tune the planning heuristic. There are six experimental parameters which are introduced throughout this section whenever they become relevant: Travel Costs Approximation (CFA) method, penalty scaling factors, new cluster costs (ncc), Acceptable Overflow Probability (AOP), installed sensors, and the length of the planning horizon. The output is a set of routes, which indicate which containers should be emptied and in what sequence they should be emptied. The three phases of the solution approach are elaborated in more detail in the remainder of this section.

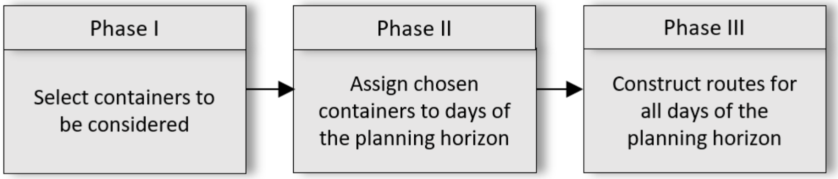

Figure 8: Decomposition of proposed solution approach

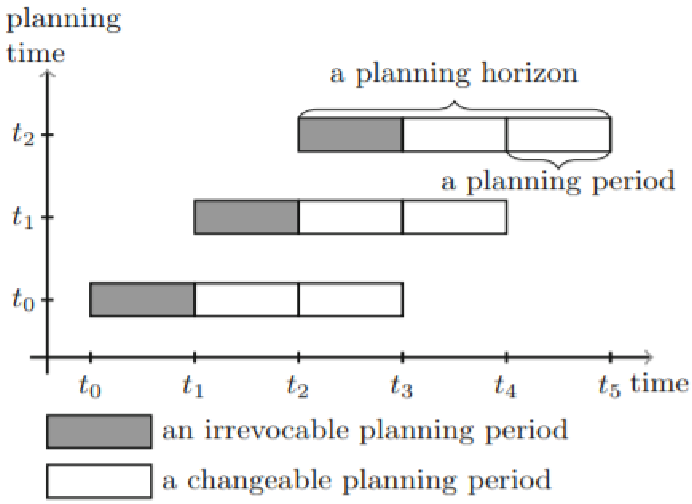

Figure 9: Example of a rolling planning horizon framework (Wang & Kopfer, 2015)

Phase I: Container selection

The first phase of the proposed solution approach selects the containers that are considered in the subsequent planning phases. Because the solution approach is intended to solve large problem instances, it can be beneficial to consider which containers to consider and which to temporarily ignore to improve the computation speed. To be able to select containers, a criterion should be formulated to decide which containers are deemed relevant and should be selected and which containers should be disregarded.

For this purpose, previous researchers use customer classifications, such as must-, may-, and no-go customers based on delivery urgency (Mes, Schutten, & Rivera, 2014) or a categorization of customers into critical, impending, or balancing customers (Campbell & Savelsbergh, 2004). Another approach is proposed by Bard et al. (1998), for each customer they determine their optimal delivery day. If this day is within a pre-determined planning horizon, the customer is selected. The optimal delivery day is found by calculating the optimal replenishment interval between two visits by minimizing the total expected delivery costs, including the penalty costs occurred on customer stockout. In a companion paper to Bard et al. (1998), Jaillet et al. (2002) argue that there must exist an optimal policy that schedules the next customer visit after a constant optimal replenishment interval. This argument is based on the assumption that the demand follows a known stationary stochastic process, which we also assume. Therefore, we are dealing with a renewal stochastic process. Each time a container is emptied marks a renewal point, at which the renewal process restarts.

Our proposed solution approach uses a concept similar to that of Bard et al. (1998). However, instead of an optimal delivery day, we introduce a new variable: the desired emptying day (DED_i). If the DED_i is within the planning horizon, container i is considered relevant and is selected to be considered in the subsequent phases. Containers that are not selected are disregarded during the planning cycle of this day, they are eligible for selection again the next day.

The DED_i is also used in the second phase of the proposed solution approach and serves to guide the timing decisions which determines on which day a container is emptied. Therefore, the DED_i should be determined in such a way that containers are not emptied too late, nor too early. If waste deposits are deterministic, this task would be trivial: the day before the container overflows would be the DED_i. However, because deposits are stochastic, using expected deposit volumes is insufficient. Therefore, we propose a more refined approach that uses a threshold: the acceptable overflow probability (AOP). The AOP is the risk, that the decision maker considers acceptable, that a container overflows before it is emptied. The DED_i can then be determined by finding the day before the overflow probability ($OP_{i,t}$) exceeds the AOP. The fill level of a container after x days can be modeled using the probability density function of the Gamma distribution, as can be seen in an example of a random container in Figure 10.

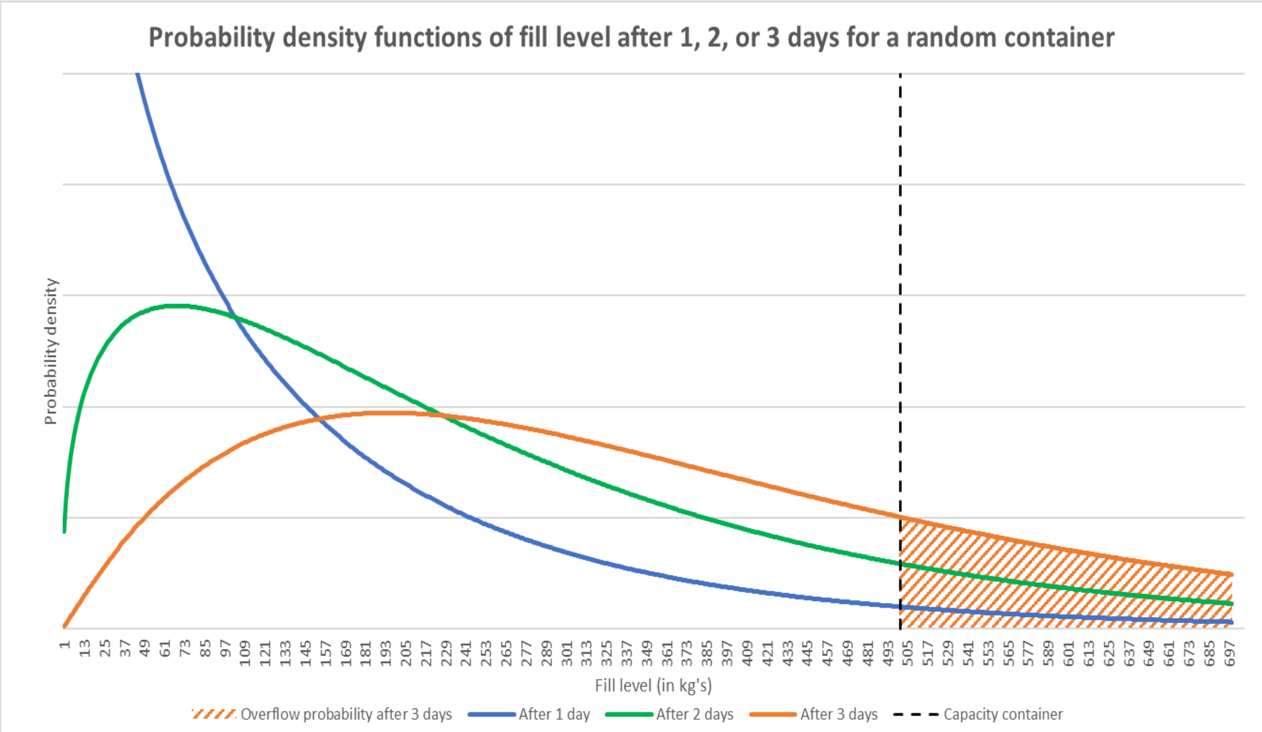

Figure 10: Probability density functions of fill level after 1, 2, or 3 days for a random container

By using the AOP to determine the DED_i, instead of more straightforward measures such as the expected day at which a container is full, we account for the stochasticity of the demand. Moreover, it gives the possibility for a conscious decision concerning the trade-off between emptying too late or too early. A risk averse approach, adopted by choosing a low AOP, will result in fewer overflowed containers. However, it also leads to many cases of containers being emptied too early with low fill levels. Choosing a high AOP has the opposite effects.

Both, the AOP and the length of the planning horizon, are used as parameters during the numerical experiments to study their effects and find appropriate values for both. Algorithm 1 shows the implementation of the container selection algorithm of Phase I. The result of the algorithm is a set of all selected containers which are used as input for the next phase.

Algorithm 1: Container selection algorithm

Algorithm 1: Container selection algorithm

```
Input:        set of all containers (C), acceptable overflow probability (AOP), length of planning horizon (T)
Output:       set of selected containers (SC)
Parameters:   desired emptying day (DEDi)

Initialize, SC = ∅
For each container i in set C
|     Calculate DEDi based on AOP and demand ~ Gamma (ki, θi)
|     If DEDi ≤ T then
|     |     add Ci to SC
|     End if
End for
Return SC
```

Phase II: Day assignment

During the second phase of the proposed solution approach, all containers selected in Phase I are assigned to days in the planning horizon, which we call the day assignment. Once all containers are assigned to days, the routes can be constructed within each day, which is done in Phase III. The goal of this phase is to decide which containers should be emptied at what time, i.e., taking the timing decision. As discussed in Section 3.1.4, the timing decision should consider both time and space dimensions simultaneously because of the interconnectedness between the two dimensions. To do this, containers are combined that are both adjacent and have similar DED_i's. By considering both the location and timing aspects simultaneously, the time and space dimensions are considered in the decision simultaneously.

The decomposition scheme of Bard et al. (1998) also contains a step in which customers are assigned to days of the planning horizon. Their day assignment is primarily based on the optimal delivery day, as explained in the previous section. However, customers are not necessarily scheduled on their optimal day. One of the goals of Bard et al. (1998) is to balance customers over the days of the planning horizon, therefore it may be necessary to move customers to other days. This is possible, however, as it is suboptimal, incurs a penalty costs, which Jaillet et al. (2002) introduce as the incremental costs. Bard et al. (1998) use a generalized assignment problem that balances the workload over the days of the planning horizon while also minimizing the total incremental costs. However, the day assignment of Bard et al. (1998), primarily based on the incremental costs, does not sufficiently account for the interconnectedness of the time and space dimensions. The incremental costs as introduced by Jaillet et al. (2002) are calculated for each customer independently. This means that the distances between customers are not considered during the day assignment, while this has serious implications on the resulting routing possibilities.

The proposed solution approach attempts to make the day assignment considering both the time and space dimensions of the IRP. In this way, the timing decision can consider spatial aspects of the IRP, such as the locations of containers. This then enables us to make the decision to shift containers from their DED_i to other days where other adjacent containers are scheduled to accommodate more efficient routes. The goal of this phase is to combine containers that are adjacent and have similar DED_i's into clusters. Clusters are collections of, preferably nearby, containers. Each day of the planning horizon can contain multiple clusters, but a cluster is assigned to one specific day. All containers in a cluster are emptied on the day to which the cluster is assigned.

Containers should be assigned to clusters which are both close or on their DED_i and contain adjacent containers. To judge which clusters is appropriate for each container, we introduce two variables, each accounting for one dimension of the IRP: the cluster fitness approximation (space) and timing penalty costs (time). To assign containers to clusters, we adopt an approach similar to the cheapest insertion heuristic commonly used in VRPs. Instead of inserting customers into routes, as is its function in the VRP, containers are inserted into clusters. The cheapest insertion for a container is found by iterating over all clusters and calculating the insertion costs $(ic_{i,j})$ of each insertion option, as is shown in Equation 1. The calculation methods for both the cluster fitness approximation $(CFA_{i,j})$ and the timing penalty $(TP_{i,t})$ are elaborated later this section.

$$ic_{i,j} = CFA_{i,j} + TP_{i,t} \tag{1}$$

In addition to the possibility of adding a container to an existing cluster, it is also possible to create a new cluster. New clusters can be created when the overall cheapest insertion exceeds a certain threshold parameter called the new cluster costs (ncc). The ncc is an experimental parameters which is studied during the numerical experiments of Chapter 6. When a new cluster is created, the initiating container is first added as the sole container and the cluster is assigned to the DED_i of that initial container.

At the start of each planning cycle, Algorithm 2 creates three clusters on each day of the planning horizon. This is done to ensure that the seed container, i.e., the first container considered for insertion, is not always the first container for which a cluster is created. Without the free clusters on each day, a seed bias caused by far most clusters to be created on the first day of the planning horizon causing an imbalance in the planning.

The exact implementation of the day assignment algorithm is shown in Algorithm 2. Similar to a cheapest insertion heuristic, all possible heuristics are evaluated and the cheapest is performed, restarting the procedure. Before evaluating an insertion, the feasibility is checked. Insertions can be deemed infeasible if clusters exceed the permitted waste load, which is based on the waste storage capacity of collection vehicles. Moreover, Algorithm 2 contains one supplementary algorithm that can be found in Appendix 5. Algorithm 2.1 is used to schedule containers that overflowed the day before. As the emptying of these containers is mandatory, it is beneficial to take their locations into account when scheduling the remaining containers. The overflowed containers are clustered together if they are close enough and create a new cluster if not. Step 3 of the algorithm is used to create and initialize three free clusters, this is discussed in more detail together with the costs for adding new clusters later this section.

Because a lot of individual insertions are considered during the execution of this algorithm, calculating the costs of each insertion should be fast. Therefore, we introduce several cost approximation methods that are used to estimate the potential impact of decisions. There are two cost measures in Algorithm 2 that should be approximated: the CFA and the timing penalty costs. The cost approximation methods for these cost measures are discussed in the following sections.

Algorithm 2: Day assignment algorithm

Algorithm 2: Day assignment algorithm

Input: set of selected containers (SC) and set of overflowed containers (OC)
Output: set of clusters (CL)
Parameters: set of unassigned containers (UC), insertion costs of container i to cluster j ($ic_{i,j}$), costs of creating a new cluster (ncc), cheapest possible insertion for container i (CI_i), overall cheapest insertion (OCI).

```
Initialize, UC = SC,  CL = ∅
Cluster overflowed containers of the previous day using Algorithm 2.1 (Appendix 5)
Create and initialize three free clusters on each day of the planning horizon
While  UC ≠ ∅  do
|     re-initialize, OCI = 999
|     For  each container i in UC
|     |     For  each cluster j in CL
|     |     |     check cluster restrictions
|     |     |     calculate ic_i,j using Equation 1
|     |     |     save cheapest ic_i,j
|     |     End for
|     |     CI_i = min{cheapest ic_i,j ; ncc}
|     |     If  CI_i < OCI then
|     |     |     OCI = CI_i
|     |     End if
|     End for
|     perform insertion associated with OCI
|     remove inserted container from UC
End while
Return  CL
```

Timing penalty costs approximation

The timing penalty costs represent the additional costs that are incurred because a container is not emptied on its DED_i. The timing penalty costs are divided into three components, as is shown in Equation 2: the penalty for emptying too late ($PEtL_{i,t}$), on time, and too early ($PEtE_{i,t}$). Both penalties are preceded by a penalty scaling factor (f_{tl} and f_{te}), these factors are used during the numerical experiments to counter possible biases and evaluate the trade-off between emptying earlier or later.

$$TP_{i,t} = \begin{cases} f_{tl} * PEtL_{i,t} & , \quad t > DED_i \\ 0 & , \quad t = DED_i \\ f_{te} * PEtE_{i,t} & , \quad t < DED_i \end{cases} \tag{2}$$

Emptying a container later than its DED_i implies that the risk of overflowing for that container, $OP_{i,t}$, now exceeds the AOP. If a container overflows before it is emptied, the municipality is notified and obligated to empty that container the next day. Because we want to prevent this from happening more often than is considered acceptable as indicated by the AOP, the $PEtL_{i,t}$, as shown in Equation 3, is based on how much the AOP is exceeded in combination with the costs of individually emptying that container. The costs of individually emptying a container are denoted by the single container routing costs ($SCRC_i$), which is the worst possible scenario that would only occur if no other containers are emptied that day.

$$PEtL_{i,t} = \left[OP_{i,t} - AOP\right] * SCRC_i \tag{3}$$

The $OP_{i,t}$, see Equation 4, is calculated using the cumulative distribution function, which calculates the probability that, at time t, the total waste volume in container i exceeds its capacity. The $SCRC_i$, see Equation 5, are based on an individual route from the wharf to the container, to either the waste processor or satellite facility, and back to the wharf.

$$OP_{i,t} = Prob\left(u_{i,t} > w_i\right) \tag{4}$$

$$SCRC_i = c_{wh,i} + \min(c_{i,wp} + c_{wp,wh}\ ;\ c_{i,sf} + c_{sf,wh}) \tag{5}$$

Contrastingly, there are also costs associated with emptying a container too early. These costs are mainly because the earlier a container is emptied, the more frequently it is emptied per time period. This is because of the renewal stochastic process described earlier. If a container is emptied earlier than necessary, the renewal stochastic process is also renewed earlier. As the optimal replenishment interval remains constant, this means that the container is emptied more often than under the optimal visiting policy. The $PEtE_{i,t}$, see Equation 6, is therefore based on the fraction of additional visits that are required within the expected interval length (EIL_i). Where the EIL_i is the time between two consecutive $DED_i's$. The number of extra visits is then multiplied with the cluster fitness approximation costs ($CFA_{i,j}$) of emptying that container. The $CFA_{i,j}$ are the costs incurred by adding container i to route or cluster j. These cluster fitness costs are based on the chosen method of travel cost approximation which are elaborated in the next section.

$$PEtE_{i,t} = \left(\frac{DED_i - t}{EIL_i}\right) * CFA_{i,j} \tag{6}$$

Travel costs approximations

Because of the desired speed of the calculations for the insertion costs, it is not desirable to solve the entire VRP or VRPSF to evaluate the travel costs for every possible insertion. However, we still want to estimate how good a container fits into a cluster. For this purpose, we formulate four CFA methods that use different methods to estimate the fitness of containers to be added a certain cluster. The four methods are called: Cluster Centroid, Cluster Proximity, Cheapest Insertion, and Daganzo, named after a VRP distance approximation heuristic by Robusté, Daganzo, and Souleyrette (1990). The methods are elaborated in Table 8 and tested during the numerical experiments in Chapter 6.

Table 8: Description of the different Cluster Fitness Approximation (CFA) methods

CFA method	Description
Cluster Centroid (CC)	Measures the Euclidean distance between the candidate container and the cluster centroid.
Cluster Proximity (CP)	Measures the Euclidean distance between the candidate container and the closest container already assigned to the cluster. This distance is then doubled.
Cheapest Insertion (CI)	Measures the marginal costs of inserting the candidate container in a route between the two closest containers already assigned to the cluster, temporarily called A and B. The marginal costs are calculated by comparing the routing costs of the route $[A \rightarrow \mathrm{B}]$ with those of route $[A \rightarrow \mathrm{candidate} \rightarrow \mathrm{B}]$.

Daganzo approximation	Calculates the marginal costs of adding the candidate container in a similar way as in the CI method. However, in this method, the routing costs are approximated using the VRP approximation of Robusté, Daganzo, and Souleyrette (1990) as given in Equation 7.

$$VRP\ distance\ approximation = \left[0.9 + \frac{kN}{C^2}\right] * \sqrt{AN} \quad (7^*)$$

** where k = area shape constant, N = number of customers, C = vehicle capacity, and A = area size*

Phase III: Route construction

The third phase of the proposed solution approach converts the day assignment made in the previous phase to feasible routes. The result of the day assignment is a set of containers that are assigned to clusters within a day. However, as the clusters are of widely varying sizes, they are not suitable to construct routes with. Therefore, we consider the input of this phase as a set of containers. During this phase, routes are only constructed for the first day of the planning horizon, this is also the only day of the planning that is actually executed, as the entire algorithm will run again the next day as dictated by our rolling horizon framework.

The resulting problem is a VRPSF, which we solve in four steps: (1) preprocessing the container set to make clusters suitable for tour construction using an adapted k-means algorithm, (2) applying a nearest insertion heuristic within each cluster to find a feasible sequencing, (3) string tours together into feasible routes, and (4) improve the current solution using a 2-opt algorithm. These four steps are elaborated further in the following sections.

Preprocessing using a k-means algorithm

The first step in our approach is to partition all containers in the container set into clusters that can be emptied in one tour. We define a tour as a routing sequence which the collection vehicle starts while empty and ends at a disposal facility, where it is once again emptied. Between this empty starting point and the end at the disposal facility, the tour consists of a sequence of containers that should be emptied. In contrast with a route, with starts and ends at a wharf, a collection vehicle can perform multiple tours per day. Moreover, a route consists of one or multiple tours. A cluster suitable for tour construction should therefore comply with the capacity restrictions of the collection vehicle.

Because we want to minimize the number of detours and processing times at disposal facilities, we want to minimize the number of tours and consequently clusters. This is done using a k-means algorithm (Geetha, Poonthalir, & Vanathi, 2009), which function is to partition a set of containers into k clusters based on the shortest path to the cluster centroid and a known cluster capacity. The implementation of the algorithm can be found in Appendix 5. Before the k-means algorithm can be started, the number of desired clusters (k) is determined using Equation 8. The algorithm then chooses k starting points and assigns all containers to their closest cluster. However, as tours and thus clusters have a maximum capacity, this is done in order of priority which is determined using Equation 9. The container with the highest priority gets assigned first, then the second highest, and so forth. If a container cannot be assigned to its closest cluster, for example due to capacity restrictions, its priority is recalculated using its second closest cluster, after which all remaining containers are reconsidered. After all containers are assigned, the centroids of all k clusters is recalculated and the container assignment procedure is repeated. This is done iteratively for ten iterations.

$$k = \left\lceil \frac{total\ load\ of\ all\ containers}{(1 - vehicle\ buffer) * vehicle\ capacity} \right\rceil \tag{8}$$

$$priority_i = \frac{expected\ fill\ level\ container}{distance\ to\ closest\ cluster} \tag{9}$$

Tour construction using a nearest insertion heuristic

After the first step, a number of clusters is created which are to be converted into tours during this step. This is done using a straightforward nearest insertion heuristic. Tours are constructed by starting from the wharf and performing the nearest insertion, which is evaluated by calculating the difference between the travel distance before and after adding a container to the tour. All possible insertions are considered and the cheapest is performed in an iterative manner until all containers are included in the tour. When all containers are added, a disposal facility is chosen. This decision is based on the travel costs between the last container, the disposal facility, and the wharf, and a satellite facility fee if the satellite facility is used. More information on the fee associated with using the satellite facility can be found in Appendix 6. The cost of both options are considered and the cheapest disposal facility is chosen. This procedure is performed for all clusters created during the previous step, resulting in a set of tours.

Creating routes from tours

As the tours created in the previous step are seldom long enough to fill an entire workday and it is undesirable to visit the wharf each time a new tour is started, this step attempts to combine tours to create complete routes.

As we want to execute the planning using as little collection vehicles as possible, we attempt to minimize the number of routes. This is done by starting with an empty route and attempting to fill it with the longest tour. After the longest tour is added, the second longest tour is added, and so forth. When a tour cannot be added to a route anymore a new route is created. In this way, the number of routes and thus required collection vehicles is minimized.

When all tours are assigned to routes, the tours should still be sequenced within the route. The tours are stripped of their wharf entries and are added to routes using a nearest insertion heuristic considering only the beginning and endpoints of each tour. This result in routes similar to those shown in Figure 11.

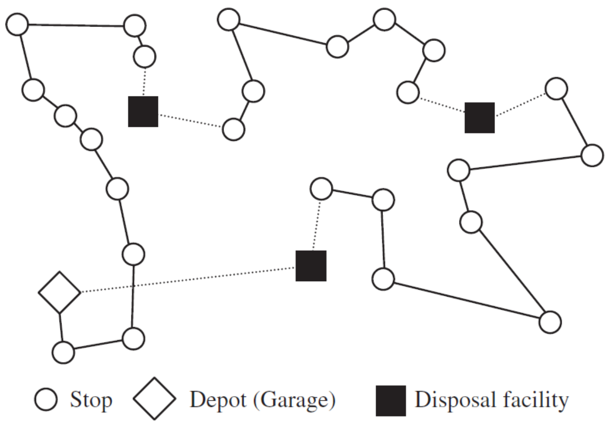

Figure 11: Collection route with multiple disposal facilities (Kim, Kim, & Sahoo, 2006)

Applying a 2-opt improvement algorithm

As a last step, when feasible routes are already constructed, the routes are improved using a 2-opt algorithm. The 2-opt algorithm starts by removing two edges from the route and reconnects the resulting paths with two different edges (Nilsson, 2003), as shown in Figure 12. This is called a 2-opt move, the move is only performed if the resulting route is shorter, otherwise the original edges are restored. 2-opt moves are attempted for each combination of edges until no improvement is possible anymore. The precise implementation can be found in Appendix 5.

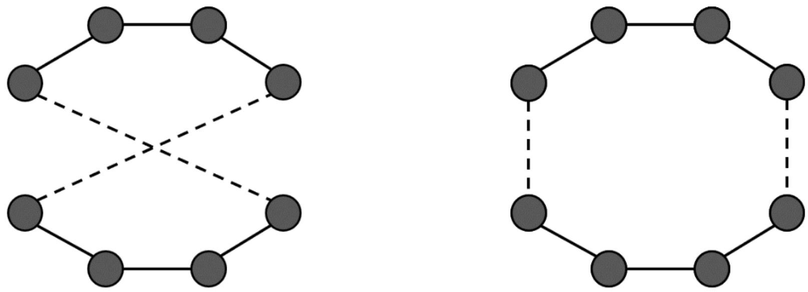

Figure 12: A 2-opt move

4.4 - Conclusion

This chapter starts by describing the problem, after which a novel solution approach is proposed. In doing so, answering research question 3: How can a novel planning methodology be designed for the waste collection in Amsterdam?

This chapter starts by giving a description of Amsterdam's waste collection problem, including several notations which support the development and elaboration of the solution approach. From this formulation, two decisions become apparent: which container should be emptied at what time and what route to use to visit the chosen containers. Moreover, the circumstances surrounding the waste collection system of Amsterdam are elaborated. As noted in Chapter 3, the contribution of this thesis is to propose a novel solution approach that simultaneously considers the time and space dimensions of the IRP.

The proposed solution approach consists of three phases: container selection, day assignment, and route construction. In the first phase, containers are selected that are considered relevant during the planning horizon. These containers are used as input for the subsequent phases, the other containers are ignored to reduce the computation time. The second phase assigns all containers to days of the planning horizon. This is done considering the effects of both the time and space dimensions of the IRP. Once it is clear which containers are to be emptied on which day, the third phase uses routing heuristics to construct final routes. During each phase, several experimental factors are described that are used to calibrate the proposed solution approach to the characteristics of Amsterdam. The effects of these factors are studied in the following chapter.

Chapter 5 - Simulation model

This chapter starts by describing the developed simulation model that is used to evaluate the different planning methodologies, while also discussing some simplifications and assumptions (Section 5.1). After that, the simulation model is verified and validated using several techniques (Section 5.2). Subsequently, the design of experiments is explained (Section 5.3) and the applied replication/deletion approach is elaborated (Section 5.4). Lastly, this chapter is concluded by answering the research questions associated with this chapter (Section 5.5).

5.1 - Description of the simulation model

This section describes the simulation model that is used to evaluate different planning methodologies and the effects of the experimental parameters that are listed in Section 5.3. Firstly, the situation which is subject to the simulation study is described in more detail, including the demarcation of the scope relative to the situation described in Chapter 2. After that, the general structure of the simulation model is given.

The purpose of the simulation model is to give insight into the effects of the experimental factors and evaluate different planning methodologies. This is done by mimicking the waste collection system of Amsterdam, changing certain aspects and noting the differences in the planning performance. To reduce the problem size, which improves computation speed and makes the effects of different planning decisions easier to comprehend, the scope of the simulation model is narrowed in two ways. Firstly, by only considering one waste fraction: household waste. As our proposed solution approach assumes a homogeneous fleet, the collection of different waste fractions can be seen as separate individual planning problems. The household waste fraction is chosen as it is the largest waste fraction, incurring the most costs. Secondly, the scope is narrowed geographically by selecting one district of Amsterdam to model. For this purpose, “Amsterdam Zuidoost” is chosen which is located in the southeast of Amsterdam. This district is chosen because of its secluded location in relation to the rest of Amsterdam’s districts and the presence of both a satellite facility and waste processor. It should be noted that the waste processor located in Amsterdam Zuidoost is, in reality, not used to process household waste. However, this simplification is made to keep the decision between using the waste processor or a satellite facility relevant. Otherwise, the routing problem faced in Phase III of the proposed solution would be reduced to a VRP instead of a VRPSF, which is the actual problem facing the municipality of Amsterdam. A full list of simplifications can be found in Appendix 6.

The waste collection system as modeled in the simulation model consists of 353 underground containers that cumulatively collect an average of 11.274 tons of waste per year. The containers are scattered over an area the size of 21.7 km^2, which also contains one wharf, one satellite facility, and one waste processor. A visual representation, including an example of a route, of Amsterdam Zuidoost and the previously mentioned locations can be found in Figure 13.

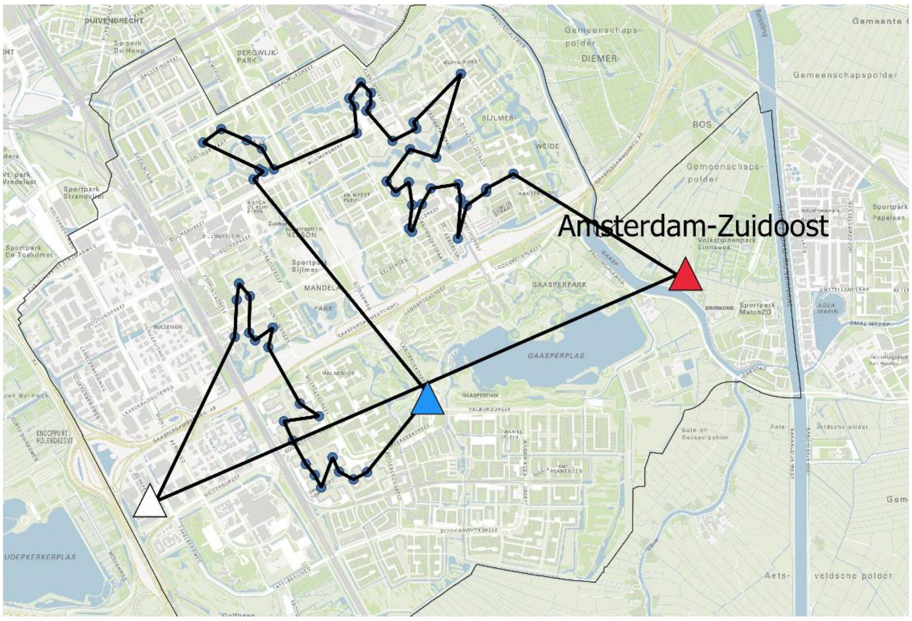

Figure 13: Visualization of routes

The simulation model is developed as a discrete-event simulation (DES). This means that time in the simulation advances by jumping to the next scheduled event (Law, 2015). In the time between two events, the state of the system does not change. The structure of the simulation model can be seen in Figure 14. One of the most important aspects of DES is the event controller, which dictates the order in which events happen and is the backbone of the simulation. The event controller governs the actions of the two actors in system: the inhabitants of Amsterdam and the waste collection department. There are two events triggered by the event controller which change the state of the system:

- The depositing of waste into containers by inhabitants,
- The emptying of containers by the waste department.

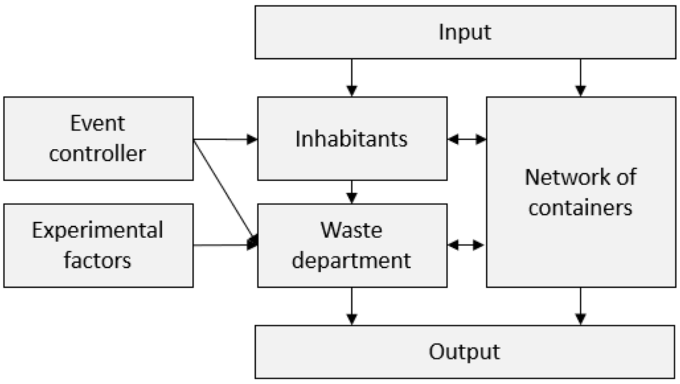

Figure 14: Structure of simulation model

The simulation model has two types of input: general characteristics of the waste collection system and aspects relating to the planning methodologies, such as the experimental factors. The general characteristics of the system include the locations of components of the logistical chain, waste disposal rates, and available vehicles. The experimental factors influence the decision making processes at the waste department, as elaborated in Chapter 4. The output of the simulation model includes the previously discussed KPIs, as discussed in Chapter 2, with several additions such as the number of routes, vehicles used, total waste collected, and computation time.

The simulation model is implemented using the Java programming language and does not contain any visualization. However, it does offer the functionality to output a plain text file that, using well-known text (WKT) formats, allowing users to export the planning to a geographical information system (GIS), such as QGIS or PostGIS. An example of the output is shown in Figure 13.

5.2 - Verification and validation

One of the challenges encountered while using mathematical modeling approaches, such as simulation, is to ensure that the model gives an accurate representation of reality. To confirm that this is the case, Law (2015) suggests using two techniques: verification and validation. Verification confirms that all desired aspects of the actual system are correctly translated into the computer program representing the simulation model. This should be done in the appropriate, chosen level of detail and according to a written list of assumptions (Law, 2015). During model validation, it is studied if the simulation model is an accurate representation of the actual system (Law, 2015). Because it is often necessary for simulation studies to make assumptions and use simplifications, simulation models only offer an approximation of the actual system. If a simulation model is an adequate approximation, it is useful during the decision-making process. To verify and validate the simulation model described in the previous section, several steps are taken as suggested by Law (2015): structured walkthroughs, quantitative component validation, and visual animation checks.

Several structured walkthroughs are performed, each focusing on different aspects of the simulation model. A structured walkthrough encompasses the following the step-by-step progress of the simulation model through time, validating at each step if the correct actions are completed. This is done both for an entire day and by following a select group of containers over a longer period of time. From these walkthroughs, no irregularities are found and the simulation model behaves as intended.

Quantitative validation is done on the waste deposits rates per container and overall waste generated. The input parameters for each individual container are, among other things, the statistical distribution modeling the amount of waste deposited in that container each day. From this statistical distribution, random numbers are used to generate random daily deposits. To validate the accuracy of this distribution, the actual observed mean and variance are compared with that of a large number of randomly generate deposits. Moreover, the entire waste generation of the district of Zuidoost for a year is compared to the amount generated during the simulation runs. Over three replications, the average waste generation during the simulation is 3.101 tons of waste, while the observed waste generation in the data collection period in Amsterdam was 3.089 tons, which is a difference of 0,4%. Moreover, in reality, the average standard deviation of the waste deposits per container is 97,7, compared to 96,9 in the simulation, a difference of 0,8%. These differences are deemed to be acceptably small.

Lastly, visual checks are performed with the use of animation. The output of each phase of the solution approach is visualized on a map of Amsterdam and the choices of the heuristic are checked for consistency and correctness. The decisions taken by the heuristic comply with what is to be expected and decisions are made on a consistent and substantiated basis.

5.3 - Design of experiments

This section describes the design of experiments by listing the experimental factors and elaborating the approach of our factorial design. Each experimental factor is a controlled variable parameter within the planning heuristic which can take on several values. All experimental factors and their possible values are shown in Table 9. The aim of the numerical experiments is to study the effects of the separate factors, but also their interaction. However, to prevent an infeasible number of experiments, not all possible interactions are studied. We devise three experiments which combine experimental factors that are expected to have a relevant interaction. This approach is similar to the one-factor-at-a-time (OFAT) approach described by Law (2015). However, unlike the OFAT approach, this approach does attempt to measure some parameter interaction. Table 9 shows which experimental factors are combined in the experiments which are described in the remainder of this section.

Table 9: Experimental factors

Factor #	Exp. #	Experimental factor	Values
1	1	Cluster fitness approximation method	[CC ; CP ; CI; Daganzo]
2	1	Penalty scaling factors	[0.1 ; 0.5 ; 1 ; 5 ; 10]
3	1	Costs of adding new clusters	[0.25 ; 0.50 ; 0.75 ; 1 ; 1.5 ; 2]
4	2	Acceptable overflow probability	[0.10 ; 0.15 ; 0.20 ; 0.25]
5	2	Sensors in containers	[yes ; no]
6	3	Length of planning horizon	[1 ; 2 ; 3 ; 4]

Experiment 1: Inventory routing

The first experiment focusses on the general solution of the IRP, which consists of the decision when to empty each container and the subsequent routing decisions. These are mainly influenced by the CFA method, penalty scaling factors, and the costs of adding new clusters. These factors regulate the creation of clusters as described in Phases II and III of the proposed solution approach. Because all three factors influence the main components of the timing decision, some interaction between the factors is expected.

The interaction between the CFA method and the penalty scaling factors is especially important, as they represent or directly influence the two cost factors used during the timing decision: travel costs and timing penalty costs. Because the different CFA methods use different ways to assess a container's fitness to a cluster, it is not certain that that they give a fair assessment of the cost they are supposed to approximate. The penalty scaling factors are implemented to offset any potential tendencies the current measures might have.

Experiment 2: Dealing with stochasticity

The second experiment studies how the proposed solution should handle the stochasticity of the waste collection system. The primary source of stochasticity in the system is the uncertain amount of waste that is deposited daily. Two factors are related to this stochasticity: the AOP and the potential implementation of sensors.

While using sensors in containers, a lot of uncertainty is taken away from the decision-making process. This might enable other, more extreme, values of the AOP to be feasible, potentially resulting in higher service levels than could be achieved without sensors.

Experiment 3: Length of planning horizon

The third experiment evaluates the trade-off between the solution quality and the required computation time. This is done by altering the length of the planning horizon used in Phase I of the proposed solution approach. The longer the planning horizon, the more information is considered by the planning heuristic, which is expected to result in a better solution quality. However, this comes at a cost, as considering more information slows down the heuristic, increasing the required computation time.

Standard settings for experiments

As the three experiments are performed consecutively, standard settings for the planning methodology are required to use while the results of the impending experiments are still unknown. These standard settings are used for each experimental factor until experiments have offered better alternatives. The values of the standard settings are chosen based on a series of initial exploratory experiments and are given in Table 10.

Table 10: Standard settings of the experimental factors

Factor #	Experimental factor	Possible alternatives
1	Travel costs approximation method	Cluster Centroid
2	Penalty scaling factors	1
3	Costs of adding new clusters	0,75
4	Acceptable overflow probability	0,20
5	Sensors in containers	No
6	Length of planning horizon	3

5.4 - Replication and deletion approach

This section discusses the replication/deletion approach, as described by Law (2015), to ensure valid and reliable simulation results. The replication/deletion approach involves choosing a warm-up period, or deletion period, and an appropriate number of replications.

Warm-up period

Because the modeled system has no natural end or beginning for experiments, the simulation is called a nonterminating simulation. In a nonterminating simulation, the initial conditions can have an undesired influence on the performance of the system in the early stages, called the initialization bias (Law, 2015). For example, in our simulation, all containers start completely empty. In reality, this is an improbable state of the system. To make sure the initial conditions do not affect the performance measurements, it is necessary to wait until the system reaches its steady state before any collecting performance data. The steady state is the state in which the actual system continually operates and thus the state which we are interested in monitoring. The period until the system reaches its steady state is called the warm-up period.

To determine an appropriate warm-up period, we use Welch's graphical procedure as described in Law (2015). Welch's procedure involves making n independent replications, with a large enough number of observations m. For each observation, the average over all replications is calculated and converted into a moving average with window w. The warm-up period can then be identified by finding the period after which the moving average appears to converge. This procedure is fully executed in Appendix 7. From this analysis, we determine that a warm-up period of 25 days is appropriate. Therefore, we choose the run length of each replication to be 125 days, of which the first 25 are disregarded.

Number of replications

Because the simulation model uses stochasticity, it is insufficient to perform one single replication of each experiment as this would have little statistical significance. Therefore, we should perform multiple independent replications and calculate the mean and confidence intervals of all performance measurements.

Because the simulation model contains stochasticity, it is insufficient to perform one single replication of each experiment. To obtain reliable simulation results, multiple replications should be performed. We use the sequential procedure proposed by Law (2015) to determine how many replications are required. The sequential procedure involves increasing the number of replications until Equation 10 is satisfied, where $\delta(n, \alpha)$ is the confidence interval half-length, $\bar{X}(n)$ is the point estimate for μ, and γ' is the adjusted relative error.

$$\frac{\delta(n, \alpha)}{|\bar{X}(n)|} \leq \gamma' \tag{10}$$

The complete calculations can be found in Appendix 7. From these calculations, we conclude that, with a confidence level $\alpha = 0{,}05$ and relative error $\gamma = 0{,}05$, three replications should be performed to attain statistically relevant results.

5.5 – Conclusion

This chapter describes the way the proposed solution approach is evaluated during the numerical experiments. The implemented simulation model is described, the design of experiments is discussed, and the applied replication deletion approach is elaborated. This chapter answers research question 4: How should the waste collection system of Amsterdam be modeled to allow for the evaluation of novel planning methodologies.

The proposed solution approach is tested using a simulation model of the waste collection system of the Zuidoost-district. Moreover, only containers of the household waste fraction are considered to simplify the model. This scope and level of detail is chosen as it is determined to be most suitable for the objectives of the simulation study. Using several verification and validation techniques as suggested by Law (2015), such as: structured walkthroughs, quantitative component analysis, and visual animation checks. All techniques show no irregularities or unexpected behavior, so the simulation model is accepted to be an accurate representation of the waste collection system of Amsterdam for our purposes.

Three experiments are formulated to study the effects of the experimental factors. The experiments are designed in such a way that expected interactions between factors are also evaluated. The three experiments with their respective experimental factors are shown in Table 11.

Table 11: Design of experiments

Exp. #	Experiment	Experimental factors
1	Inventory routing	Cluster fitness approximation method, penalty scaling factors, costs of adding new clusters
2	Dealing with stochasticity	Acceptable overflow probability, sensors in containers
3	Length of planning horizon	Length of planning horizon

The experiments are performed using a replication/deletion approach. This approach involves replicating the same experiment and using a warm-up period to only measure the performance in the system's steady state. In this analysis, it is determined that each experiment should consist of three replications of which the first 25 days are the warm-up period to ensure statistically relevant results. The experiments are performed during the numerical experiments of the following chapter allowing the analysis of the effects of the different experimental factors.

Chapter 6 - Numerical experiments

This chapter discusses the results of the numerical experiments described in the previous chapter. The experiments are conducted by simulating the performance of the waste collection planning of Amsterdam over a period of 100 days using the simulation model described in Chapter 5. The goal of the experiments is to find appropriate values for all experimental factors. Before discussing the numerical experiments, the hypothetical performance of Amsterdam's current planning methodology is evaluated in our simulation model (Section 6.1). This is done to ensure a fair comparison and solves the problem of missing performance indicators such as the number of overflowed containers. The first experiment concerns the experimental factors influencing the day assignment and routing phases of the proposed solution approach (Section 6.2). The second experiment explores how the stochasticity of the waste collection environment in Amsterdam should be handled (Section 6.3). Subsequently, the third experiment examines the effects of the length of the planning horizon on the planning performance and computation time (Section 6.4). Finally, a conclusion is drawn based on the results of the numerical experiments (Section 6.5).

6.1 - Current performance of Amsterdam

To put the planning performance of the proposed solution approach into perspective, it is compared to the performance of Amsterdam's current planning methodology. Unfortunately, there are several issues preventing a direct comparison between the two planning methodologies. The foremost issue is the lack of data availability of the planning performance of Amsterdam. The municipality does not gather information on one of the main KPIs: the service level. Another issue is a modification made in the simulation model in comparison to the real-life situation in Amsterdam: the ability of the waste processor to handle household waste. In reality, the waste processor in Zuidoost is solely compatible with the paper waste fraction, while in the simulation model we modify this to include household waste. This is done to increase the complexity of the problem as elaborated in Section 5.1 and Appendix 6. Because of these issues, the planning methodology of Amsterdam is mimicked in our simulation model and the resulting performance is used in the comparison to the proposed solution approach.

The current planning methodology of Amsterdam is based on fixed emptying frequencies which indicate the time intervals between emptying the container. The exact implementation and method of determining the fixed emptying frequencies are elaborated in Appendix 8. Mimicking the planning performance of the current planning methodology of Amsterdam yields the following results: 2.14 kilometers driven per collected ton of waste and a service level of 80.4%.

6.2 - Inventory routing

This section discusses the experiments in which the experimental factors that influence the day assignment and route construction phases of the proposed solution approach. The corresponding experimental factors are: the cluster fitness approximation (CFA) method, penalty scaling factors, and costs of adding new clusters. Two experiments are devised to evaluate the effects of these three experimental factors on the planning performance. Firstly, the interaction between the chosen cluster fitness approximation method and the penalty scaling factors is studied to find appropriate scaling factors for each CFA method (Section 6.2.1). Several promising configurations are selected and used in the second experiment where different costs for creating new clusters are evaluated (Section 6.2.2).

6.2.1 - Experimental factors influencing the day assignment phase

During the day assignment phase, containers are assigned to days of the planning horizon. This assignment is based on a trade-off between the expected increase in travel costs, influenced by the CFA, and the timing costs, influenced by the penalty scaling factors (f_{tl} and f_{te}). This trade-off causes an interaction between these two experimental factors, which is why they should be studied together. The influence of the penalty scaling factors is twofold: (1) negating possible biases caused by over- or underestimating the travel costs of the CFA and (2) influencing the timing decision by setting the ratio between penalizing lateness and earliness.

To evaluate the interaction between the experimental factors, we evaluate the planning performance for each possible configuration of the experimental factors. The possible CFA methods are: Daganzo, Cluster Proximity (CP), Cheapest Insertion (CI), and Cluster Centroid (CC). There are five possible levels for the penalty scaling factors: greatly reduced (0.1), reduced (0.5), normal (1), increased (5), and greatly increased (10). The results are shown in Table 12, where D*ist* shows the number of kilometers driven per ton of waste and SL denotes the service level. Results marked in boldface are identified as promising configuration which are elaborated further later this section.

Table 12: Experimental results CFA methods and penalty scaling factors

		CFA method							
PEtL	PEtE	CC		CP		CI		Daganzo	
		Dist	SL	Dist	SL	Dist	SL	Dist	SL
0.1	0.1	1,97	0,836	1,95	0,836	1,97	0,838	1,98	0,835
	0.5	1,95	0,829	1,94	0,833	1,96	0,836	1,96	0,830
	1	1,93	0,822	1,92	0,826	1,94	0,828	1,96	0,825
	5	**(1,86)**	**(0,782)**	**(1,87)**	**(0,800)**	1,90	0,800	1,89	0,798
	10	1,85	0,764	1,86	0,780	**(1,85)**	**(0,780)**	**(1,86)**	**(0,780)**
0.5	0.1	1,98	0,844	2,01	0,846	2,00	0,844	2,00	0,846
	0.5	**(1,95)**	**(0,841)**	1,98	0,843	1,98	0,844	2,00	0,838
	1	1,94	0,834	1,95	0,838	1,96	0,838	1,96	0,84
	5	1,89	0,794	1,88	0,812	1,90	0,811	1,91	0,809
	10	1,88	0,774	1,88	0,788	1,88	0,789	1,88	0,794
1	0.1	2,00	0,847	2,02	0,849	2,04	0,846	2,03	0,848
	0.5	1,96	0,843	1,99	0,846	1,98	0,843	1,99	0,846
	1	1,96	0,837	1,96	0,844	1,96	0,839	1,99	0,841
	5	1,91	0,798	1,90	0,815	1,89	0,814	1,93	0,820
	10	1,89	0,778	1,89	0,793	1,89	0,795	1,89	0,798
5	0.1	**(2,02)**	**(0,852)**	**(1,94)**	**(0,833)**	**(2,03)**	**(0,855)**	**(2,03)**	**(0,853)**
	0.5	2,00	0,847	**(2,01)**	**(0,852)**	2,01	0,848	2,03	0,850
	1	1,97	0,842	2,01	0,849	1,97	0,847	2,00	0,847
	5	1,92	0,803	1,92	0,820	**(1,92)**	**(0,827)**	**(1,92)**	**(0,825)**
	10	1,89	0,786	1,91	0,799	1,90	0,807	1,91	0,807
10	0.1	2,04	0,852	2,05	0,854	2,04	0,854	2,04	0,852
	0.5	2,00	0,848	2,02	0,852	2,01	0,852	2,03	0,848
	1	1,98	0,841	1,99	0,851	2,00	0,844	2,00	0,847
	5	1,92	0,803	1,92	0,824	1,93	0,826	1,93	0,827
	10	1,90	0,784	1,90	0,798	1,90	0,808	1,92	0,809

The results show that the penalty scaling factors have a greater impact on the planning performance than the chosen CFA method. The differences caused by different CFA methods is especially small in configurations that result in higher service levels, while it is somewhat more influential when service levels are lower, this is especially visible in the efficiency frontier shown in Figure 15. Notably, both the Daganzo and CC approximation methods are not able to make an efficient trade-off to reduce the distance driven per ton of waste, resulting in steeper declines of the service level than is seen in the other two CFA methods.

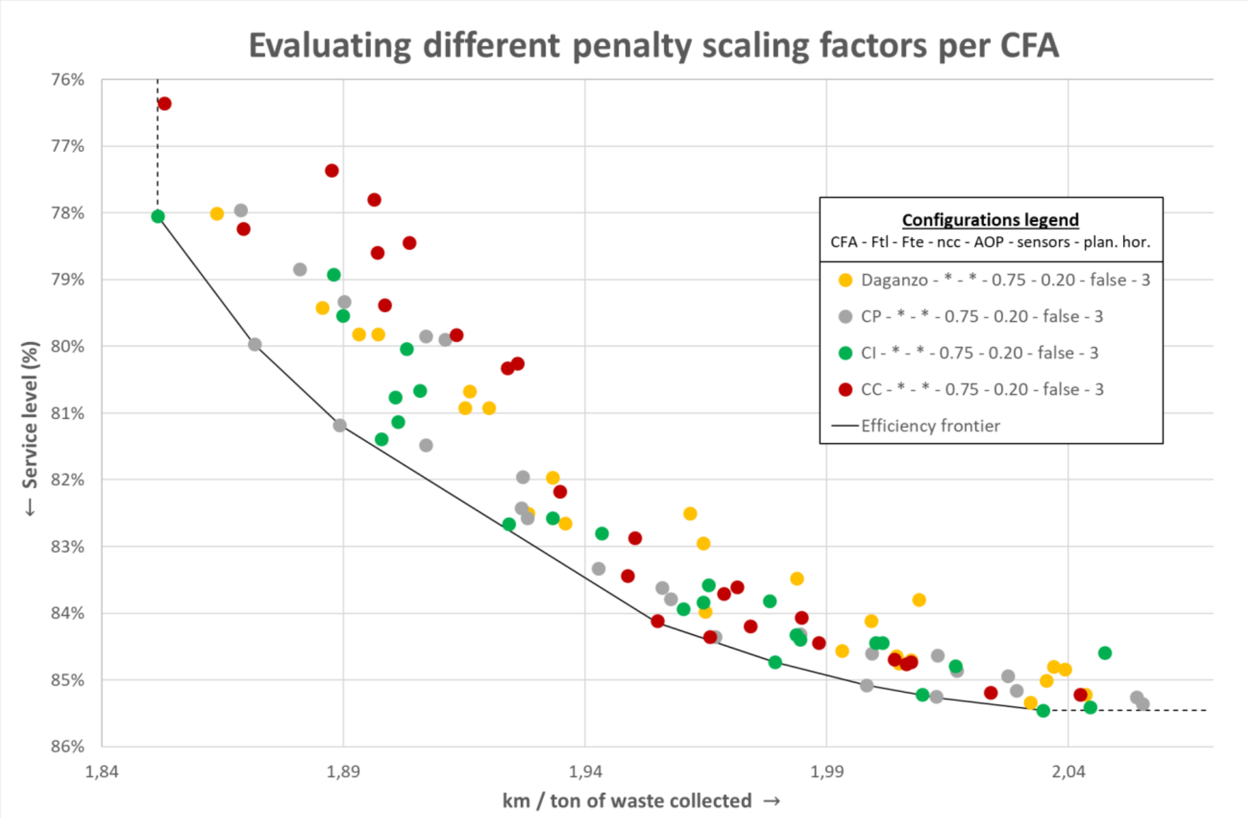

Figure 15: Evaluating the interaction between CFA methods and penalty scaling factors using an efficiency frontier

As expected, penalty scaling configurations that incentivize earlier emptying (i.e., high f_{tl} and low f_{te}) result in higher numbers of emptied containers and lower fill levels of containers upon emptying. In turn, this causes more kilometers to be driven per ton of collected waste. On the other hand, this also causes containers to be emptied earlier, reducing the overflow risk and thus increasing the service level. Contrastingly, changing the penalty scaling ratio in favor of emptying later has the opposite effects: more postponed containers, less containers emptied, lower distance per ton of waste collected, and lower service levels.

The greater influence of the CFA method for configurations where postponing is penalized less (i.e., the upper left region of Figure 15) can be explained by the dominance of the penalty for emptying later. This timing penalty is deliberately more dominant than the penalty for emptying earlier, an example of this for a random container is shown in Appendix 9. Because of this dominance, it often overrules any potential improvement in travel distance that can be achieved by emptying later. Therefore, when the f_{tl} is maintained or even increased, such as in the bottom right of Figure 15, it preserves its dominance, diminishing the influence of the CFA. When the f_{tl} is reduced, as in the top left of Figure 15, the effect of the CFA gains more significance in the timing decision, resulting in larger differences between the different CFA methods.

Conclusion and promising configurations

From Figure 15, we can conclude that no CFA method or penalty scaling configuration is clearly superior to another as the efficiency frontier consists of several different CFA methods. However, there are certainly configurations that work well together to make an efficient trade-off between the distance per ton of collected waste and the service level. These are the configurations on or close to the efficiency frontier.

Because it is not ruled out that the CFA method and penalty scaling factors have further interactions with other experimental factors, we select multiple promising configurations to continue to the next experiments with. For each CFA method, three penalty scaling configurations are chosen:

- Low distance travelled, but also low service level, called LL for 'low, low';
- High service level, but also high distance travelled, called HH for 'high, high';
- An intermediate setting that balances both KPIs, called M for 'medium'.

The promising configurations are shown in boldface in Table 12 and are summarized in Table 13.

Table 13: Defining promising configurations of Experiment 1

CFA settings	f_{tl}	f_{te}	Km / ton waste	Service level
Daganzo-HH	5	0,1	2,03	85,3 %
Daganzo-M	5	5	1,92	82,5 %
Daganzo-LL	0,1	10	1,86	78,0 %
CP-HH	5	0,1	2,01	85,2 %
CP-M	5	0,5	1,94	83,3 %
CP-LL	0,1	5	1,87	80,0 %
CI-HH	5	0,1	2,03	85,5 %
CI-M	5	5	1,92	82,7 %
CI-LL	0,1	10	1,85	78,0%
CC-HH	5	0,1	2,02	85,2 %
CC-M	0,5	0,5	1,95	84,1 %
CC-LL	0,1	5	1,86	78,2 %

Based on these first experiments, we can conclude that the relatively unmodified version of the proposed solution approach largely outperforms the current planning performance. Most configurations have a higher service level than the current performance of Amsterdam, which is 80,4%. Moreover, the current distance driven per ton of collected waste is $2{,}14$ kms, which is improved upon under each configuration of the proposed solution by 4-14%.

6.2.2 - Evaluating different costs for creating new clusters

This section studies the effect of the costs for creating new clusters, also known as the new cluster costs (ncc), on the planning performance. The ncc is expected to mainly influence the routing efficiency, as it influences the size and number of clusters that is created during the day assignment phase. If the ncc is relatively low, more smaller clusters are created in favor of less larger clusters which would arise using a higher ncc. Six different levels of the ncc are evaluated: 0.25, 0.50, 0.75, 1, 1.5, and 2. The effect of these six levels is evaluated using all promising configurations as identified in Table 13. The results of the experiment can be seen in Table 14, where the promising configurations are once again marked in boldface.

The results show that the ncc affects both the distance travelled per ton of waste and the service level. However, it is not as influential as the penalty scaling factors discussed in the previous experiment (Section 6.2.1). The results of the trade-off between the distance and service level is still highly dependent on the chosen penalty scaling factors: LL, M, and HH. This can also clearly be seen in the efficiency frontier shown in Figure 17, where the different CFA methods can be identified by color, while the different markers represent the penalty scaling settings: LL (denoted by starts), HH (denoted by triangles), and M (using spheres).

Table 14: Results of experiments studying the effects of different costs for creating new clusters

CFA method	ncc	Penalty scaling factor settings					
		Low, low		Medium		High, high	
		Dist	SL	Dist	SL	Dist	SL
CC	0,25	1,85	0,764	1,93	0,829	1,99	0,842
	0,5	1,87	0,772	1,96	0,841	2,04	0,852
	0,75	1,86	0,782	**(1,95)**	**(0,841)**	2,02	0,852
	1	1,86	0,792	1,96	0,839	2,03	0,851
	1,5	1,88	0,796	1,96	0,839	2,03	0,851
	2	1,90	0,803	1,98	0,843	2,05	0,853
CP	0,25	1,85	0,771	2,02	0,852	1,94	0,832
	0,5	1,86	0,793	2,01	0,851	1,95	0,836
	0,75	**(1,87)**	**(0,800)**	**(2,01)**	**(0,852)**	1,94	0,833
	1	1,88	0,799	2,01	0,852	1,94	0,831
	1,5	1,89	0,801	2,02	0,852	1,95	0,832
	2	1,91	0,810	2,05	0,854	1,96	0,837
CI	0,25	1,86	0,758	1,91	0,783	1,99	0,849
	0,5	1,91	0,768	1,93	0,815	2,04	0,855
	0,75	**(1,85)**	**(0,780)**	**(1,92)**	**(0,827)**	2,03	0,855
	1	1,86	0,787	1,93	0,825	2,05	0,854
	1,5	1,86	0,793	1,93	0,830	2,05	0,854
	2	1,89	0,793	1,94	0,836	**(2,09)**	**(0,859)**
Daganzo	0,25	1,87	0,758	1,91	0,795	2,03	0,848
	0,5	1,89	0,762	1,92	0,822	2,07	0,850
	0,75	1,86	0,780	1,92	0,825	2,03	0,853
	1	1,87	0,786	1,93	0,824	2,04	0,853
	1,5	1,88	0,792	1,95	0,832	2,04	0,851
	2	1,91	0,799	1,96	0,835	2,07	0,852

The results shown in Table 14 show that, in general, low ncc's result in less distance per ton of collected waste, but also in lower service levels. Higher ncc's have the opposite effect: more distance travelled per ton of waste and higher service levels. The difference made by altering the ncc is the way clusters are formed. With a low ncc, the proposed solution approach generally forms more clusters, which are also more container-dense than clusters created using higher values for the ncc. A generalized example of the created clusters under low and high ncc is shown in Figure 16.

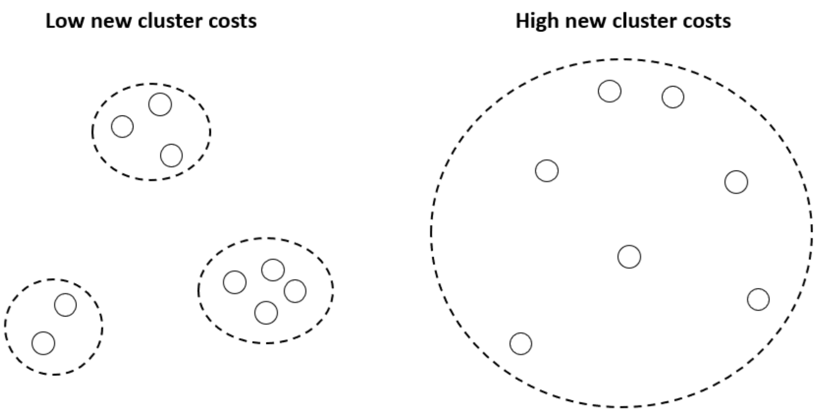

Figure 16: Average clusters created under low and high costs for creating new clusters

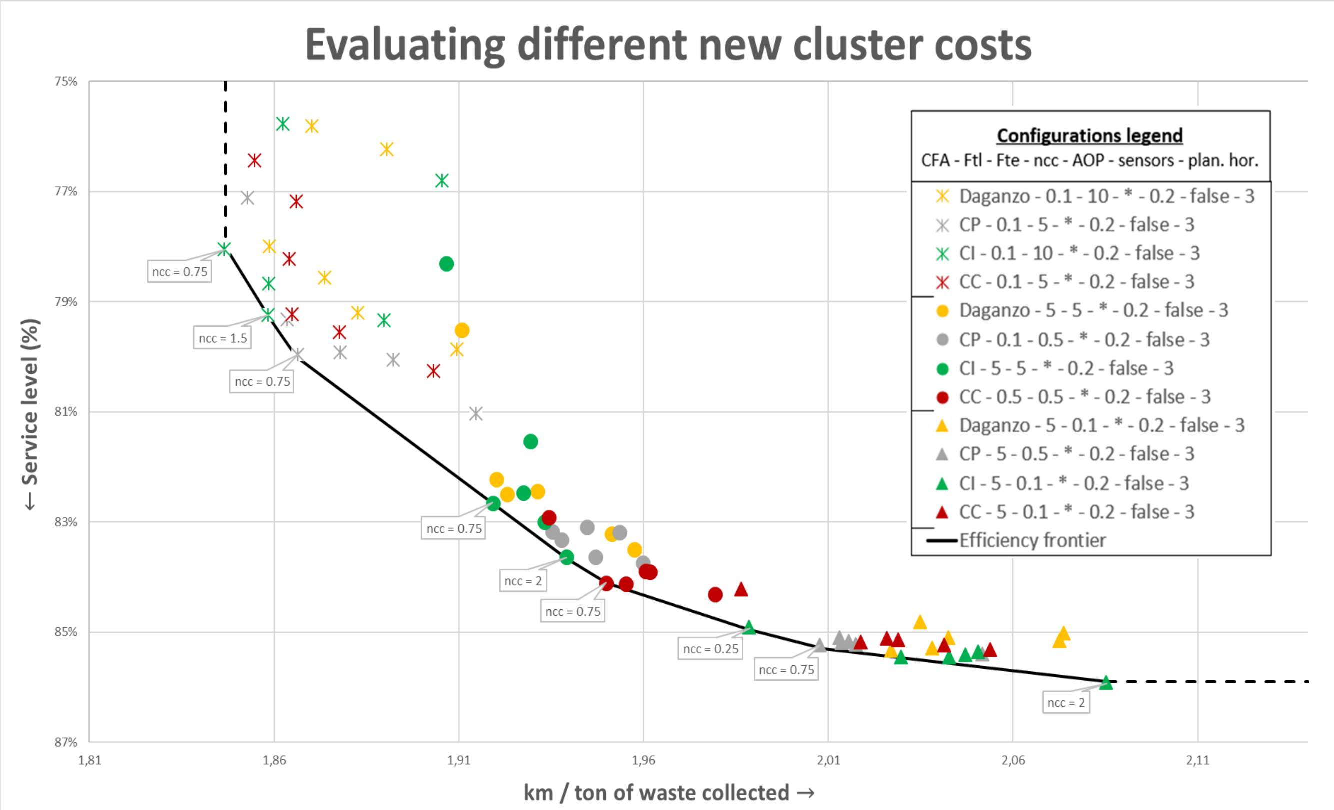

Figure 17: Evaluating different costs for creating new clusters using an efficiency frontier

The results of the experiments show that a higher number of container-dense clusters is better for the routing efficiency than clusters created with a higher ncc. This is caused by the flexibility allowed under a low ncc. As creating new clusters is easily permitted, because of the low associated costs, there is a high flexibility to create new clusters if the existing clusters are not sufficiently compatible. By adding more new clusters, new insertion possibilities present itself to all subsequent containers, strengthening the effect of the insertion flexibility. Contrastingly, if the ncc is high, new clusters are almost never created and containers are forced into clusters with which they are not actually compatible.

However, this increase in routing efficiency comes at a cost to the service level. This is caused by a side-effect of the lower ncc: as more clusters are created, they are all scheduled on the DED of the initializing container. Additionally, because of the flexibility to create new clusters easily if no compatible cluster can be found, the number of containers that is emptied on their DED increases. The result is that the lower the ncc, the more containers are scheduled on their DED. This shift is largely at the expense of emptying containers earlier. Therefore, the average container is emptied somewhat earlier when the ncc is high, which causes lower service levels.

Conclusion and promising configurations

As is shown in the results, lower values for the ncc increase the routing efficiency, expressed by the distance driven per collected ton of waste, but also decreases the service level. The opposite is true for higher ncc's. Figure 17 shows that most configurations on the efficiency frontier, i.e., the configurations that make the most efficient trade-off between the two KPIs, have a ncc of 0,75. However, to allow for extreme prioritization of one of the KPIs, we once again select several promising configurations which are marked in boldface in Table 14.

6.3 - Dealing with stochasticity

This section discusses the experiments in which we study how the stochasticity of the waste collection system of Amsterdam can best be dealt with. Section 4.3 introduced the AOP threshold which is used to deal with the stochasticity of waste deposits. The first experiment evaluates the planning performance under different levels of this AOP threshold (Section 6.3.1). Moreover, the possibility of adding sensors the underground containers to gain more information on their fill levels is explored and the potential effects are evaluated (Section 6.3.2).

6.3.1 - Evaluating different acceptable overflow probabilities

This section explores the effect of changing the AOP threshold level. The AOP is introduced as a measure to account for the stochasticity of waste deposits that is used during the day assignment phase of the proposed solution approach. It mainly affects the timing decision by determining the container's DED and influencing the formula of the timing costs. Low AOPs would be preferred by decision makers that are risk averse, as it reduces the risk that is considered acceptable. The value of the AOP threshold is differed between 10%, 15%, 20%, 25%, and 30% during the experiments. These different values for the AOP are tested on all promising configurations as identified in the previous section in Table 14. The results of the experiments can be seen in Figure 18, where the different colors indicate the promising configuration and the different points of each color represent the planning performance of that configuration with different AOP levels.

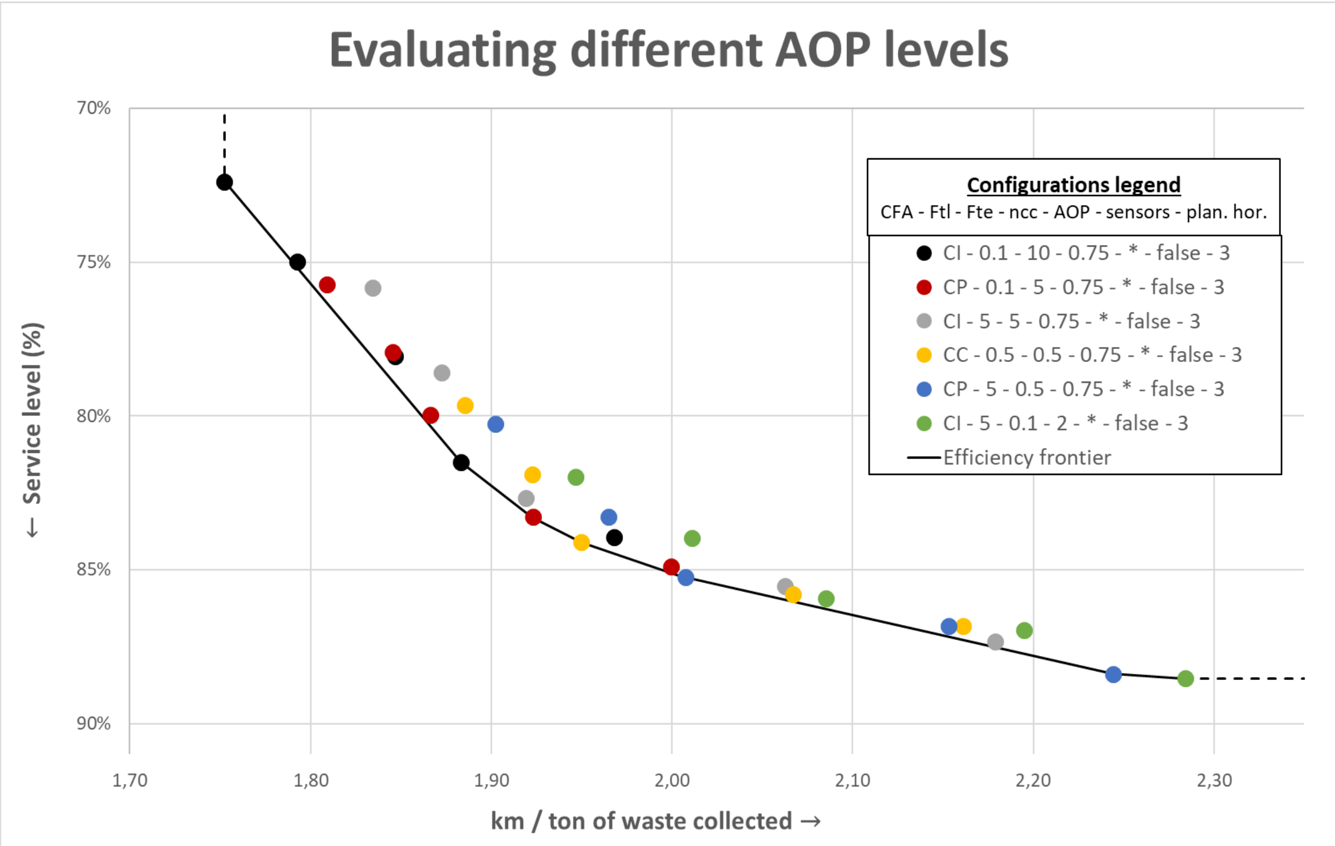

Figure 18: Evaluating different AOP threshold levels using an efficiency frontier

The results show the importance of the AOP to the trade-off between the service level and the distance driven per ton of waste. For each configuration, the difference in performance created by altering the AOP threshold can be seen in Figure 18 spanning almost half the efficiency frontier. The large impact of the AOP threshold on the performance is caused by its significant impact on the timing decision. The timing decision (i.e., emptying early, on time, or later) is largely based on a container's

DED, as this dictates when a container can be scheduled without incurring a timing penalty. While the previously studied experimental factors f_{tl} and f_{te} mainly influenced the timing decision via the timing penalty costs. The AOP directly impacts the container's DED, as well as the related EIL. The DED is defined as the day before the probability of overflowing exceeds the chosen AOP threshold level, while the EIL is the expected interval between subsequent DED's. Therefore, changing the AOP causes the starting position of the timing decision to shift by the changing DED. Moreover, the average length between emptying of a container is influenced through the EIL. The relevant trends that are observed when changing the AOP can be seen in Table 15.

Table 15: Trends when modifying the acceptable overflow probability

AOP	Average number of emptied containers	Average container fill level	Average km / ton waste	Average service level
10 %	17.994	35,8 %	2,14	86,7 %
15 %	16.225	40,0 %	2,05	85,0 %
20 %	14.562	44,6 %	1,95	82,7 %
25 %	13.362	48,5 %	1,90	80,1 %
30 %	12.062	53,1 %	1,86	77,6 %

As can be seen in Table 15, the average number of emptied containers is significantly influenced by the chosen AOP threshold. This can be explained by the previously discussed influence of the AOP on the container's EIL's. At an AOP threshold of 10% the average EIL of all containers is 3.9 days. Where, with an AOP of 30%, this increases to an EIL of 5.9 days, an increase of almost 50%.

The difference in the number of times containers are emptied is also visible in the average container fill level. Because the underlying waste deposits remain the same, the average fill level of containers drops when containers are emptied more often. When containers are emptied more often, the risk of overflowing is reduced and as a result the service level increases. However, the excessive emptying of containers also increases the number of kilometers driven per collected ton of waste.

6.3.2 - Sensors in containers

As is evident from the average fill levels of containers upon emptying in Table 15, containers are often emptied too early. A large part of this is caused by the stochasticity of the waste deposits and the desire to achieve a high service level for the inhabitants. Several papers have discussed the utilization of sensors to deal with the stochasticity of waste deposits to achieve more efficient waste collection (Johansson, 2006) (Vicentini, et al., 2009) (Mes, 2012). Amsterdam currently has not installed any sensors in their waste containers, but is considering the possibility (Municipality of Amsterdam, 2018c). To research the potential impact sensors would have on the waste collection planning performance of Amsterdam using the new proposed solution approach, we conduct new experiments where sensors are installed in each container. Before discussing the results, the exact implementation of the sensors in the simulation model and the implications for the proposed solution approach are elaborated briefly.

Implementation of fill level sensors

During normal operations of the waste collection system in Amsterdam, without sensors, the proposed solution approach uses the probability density function of container's fill levels to determine when a container should be emptied. As the only time the actual fill level of a container is known is when it is emptied, the longer that moment has passed, the more uncertain the current fill level is. However, when fill level sensors are installed in each container, the actual current fill level is known during the construction of the planning. In the simulation model used to model the waste collection system of Amsterdam, as described in Chapter 5, the planning is made each morning, the planning is carried out, and finally, the daily deposits are added to each waste container. Therefore, the decision to empty a container or postpone a container is still relevant as the container can overflow during the day if it is not emptied. Because of this, the same approach, in terms of using an AOP and DED's, is taken to planning with the addition of sensors. The difference being that the actual fill levels are updated daily which enables the DED of each container to be reconsidered each day. For example, if a deposit on the first day is less than expected, the DED moved to later in accordance with the day it exceeds the AOP.

Results of experiments

The experiments are carried out using the promising configurations of Section 6.3.1. Sensors are implemented and the AOP is differed between 5%, 10%, 15%, 20%, and 30%. Moreover, the results are compared with those achievable without installing any sensors. An excerpt of the results of the experiments can be found in Table 16, the complete table is shown in Appendix 10.

Table 16: Excerpt of experiment results, full results can be found in Appendix 10

CFA method	Penalty scaling	AOP	Sensor information: Without sensors Dist	Without sensors SL	With sensors Dist	With sensors SL	Δ (%) Dist	Δ (%) SL
CC	Medium	0,05	2,26	0,886	2,04	0,886	-9,8%	+1,4%
		0,10	2,16	0,868	1,92	0,889	-11,3%	+2,4%
		0,15	2,07	0,858	1,82	0,879	-11,8%	+2,5%
		0,20	1,95	0,841	1,79	0,861	-8,2%	+2,4%
		0,30	1,89	0,796	1,73	0,835	-8,3%	+4,9%
CP	Medium	0,05	2,37	0,905	2,15	0,911	-9,4%	+0,6%
		0,10	2,24	0,884	2,01	0,899	-10,6%	+1,7%
		0,15	2,15	0,868	1,92	0,893	-11,0%	+2,9%
		0,20	2,01	0,852	1,83	0,878	-8,7%	+3,0%
		0,30	1,90	0,803	1,76	0,848	-7,7%	+5,6%
CI	Low, low	0,05	2,09	0,862	2,02	0,881	-3,4%	+2,2%
		0,10	1,97	0,839	1,89	0,868	-4,2%	+3,4%
		0,15	1,88	0,815	1,81	0,856	-3,9%	+5,1%
		0,20	1,85	0,780	1,74	0,840	-5,7%	+7,7%
		0,30	1,75	0,724	1,69	0,809	-3,5%	+11,7%

The results show the significant improvement that can be achieved by installing fill level sensors in containers both in terms of the achievable service level and reduction of distance driven per ton of waste. The simultaneous improvement of both KPIs can be attributed to the added ability to recognize deposit trends for all individual containers supported by the sensors. By accounting for these trends, containers are less frequently emptied too early or too late.

6.3.3 - Conclusion

The results, as shown in Table 16, show the planning performance for different levels of the AOP threshold and what happens when fill level sensors are installed into all containers. Figure 19 compares the achievable efficiency frontiers in a situation where there a no sensors to one where there are sensors. The improvement in performance is clear, on average the distance driven per ton of collected waste is reduced by 8%, while the service level simultaneously increases by 4%. This are improvements on top of the previous improvements in comparison to the current planning methodology employed by the municipality of Amsterdam.

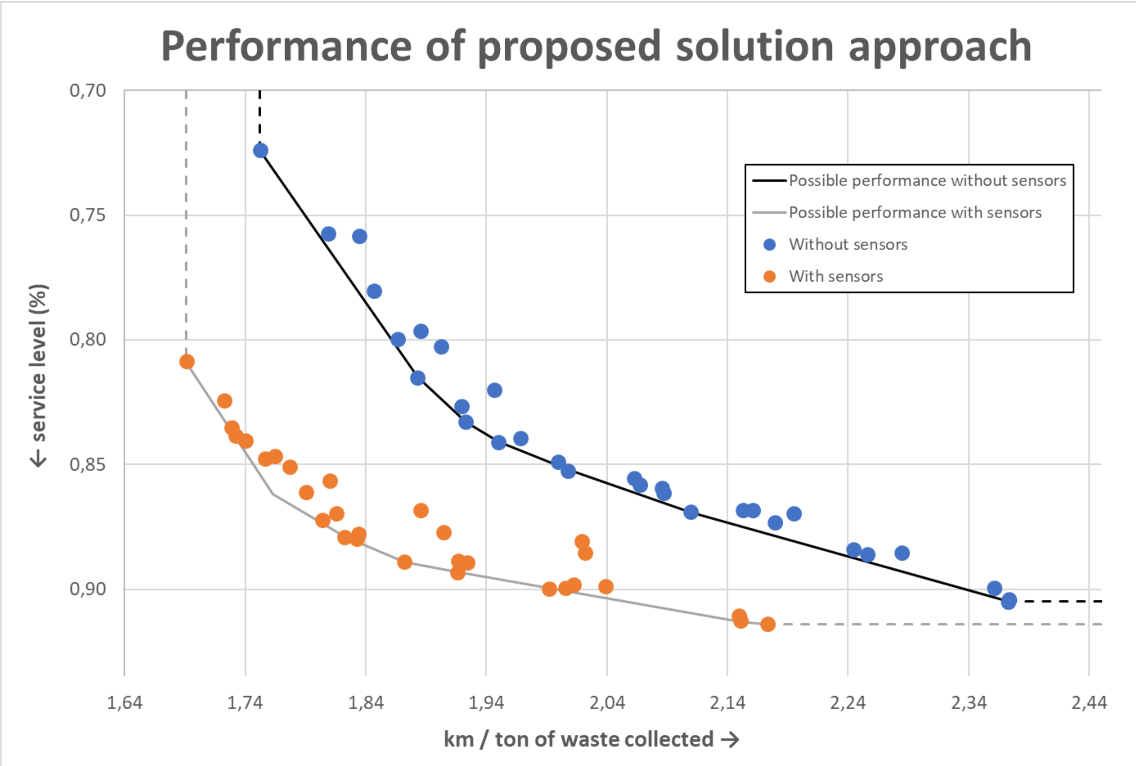

Figure 19: Evaluating the effect of installing sensors by comparing efficiency frontiers

6.4 - Planning horizon

The third experiment studies the effect the length of the planning horizon has on the planning performance. In this experiment, we do not solely consider the KPIs as used during the previous experiments, but also the computation time as this is expected to be heavily dependent on the length of the planning horizon. The computation time is important to consider because it can be a restriction in implementing the planning heuristic in real-life, where unexpected changes may require a quick recalculation of the planning. Moreover, it is also important for the scalability of the heuristic to larger problem instances. Five configurations, as identified in the legend of Figure 20, are chosen from the efficiency frontier of the stochasticity experiments of Section 6.3.2. For these configurations, the length of the planning horizon is differed between 1, 2, 3, 4, and 5 days and the results are shown in Figure 20. The labels in Figure 20 signify the length of the planning horizon used.

Adopting a longer planning horizon gives the heuristic more possibilities to schedule containers on different days than its DED, which enables the search for more efficient routes. At the same time, the timing decision is still considered by the cost function of the timing costs, of which an example can be seen in Appendix 9, which penalizes emptying days farther away from the DED heavier. However, because of the additional possibilities that the heuristic can consider, the computation time also significantly increases.

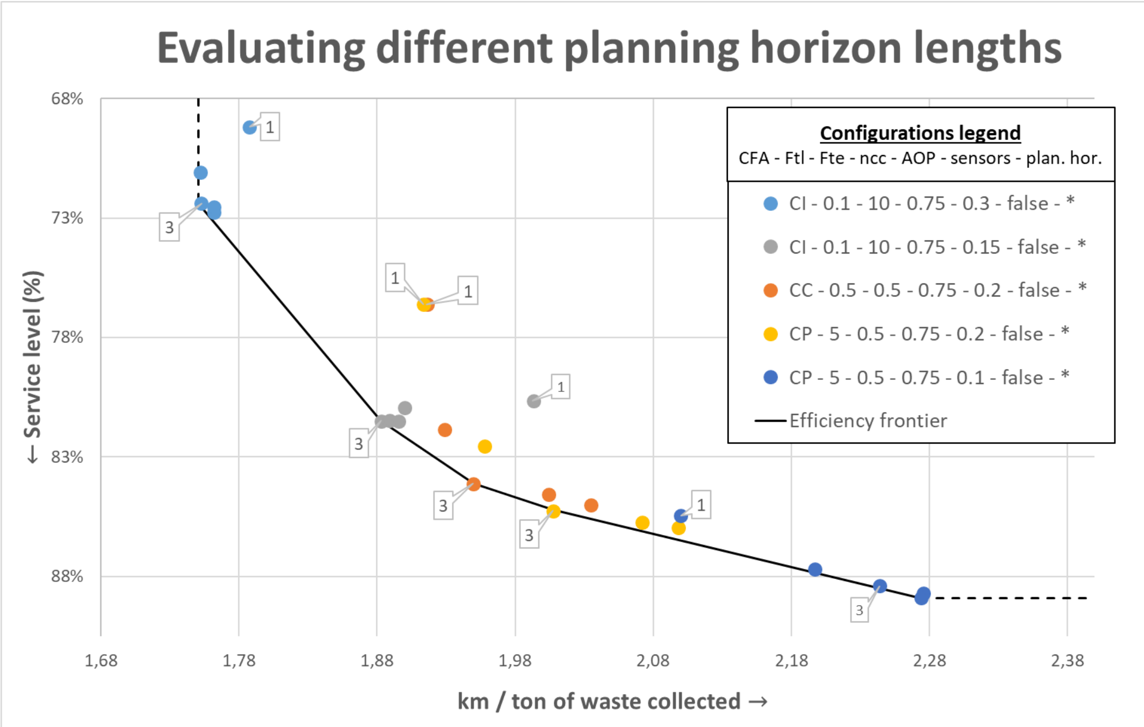

Figure 20: Evaluating different planning horizon lengths using an efficiency frontier

The results show that the configurations react differently to changing planning horizons. The most distinct difference is between the configurations that penalize earliness (blue and grey) and those that penalize lateness heavier (orange, yellow, and blue).

The former, configurations that penalize earliness heavier, are not affected as much by the length of the planning horizon. Only a planning horizon of just one day, which prohibits containers to be scheduled on a day other than their DED, performs worse on both KPIs. This is caused by two effects: the heuristic is not able to find more efficient routes by switching containers between days and the regular bias towards emptying containers earlier is not allowed. This respectively causes the routes to be less efficient, resulting in more kilometers per ton of waste, and a lower service level.

For the other three configurations, the length of the planning horizon has a greater impact on performance. These configurations have an inclination to schedule containers earlier which leads to a higher frequency of emptying containers. This results in high service levels, but also in longer distances travelled per ton of collected waste. However, when the planning horizon is reduced to one day, containers cannot be scheduled earlier, decreasing the number of emptied containers and distance per ton of waste. Although the distance per ton of waste increases with increasing the length of the planning horizon, the routing efficiency improves as is shown in Table 17. This shows that the increase in distance travelled is caused by the timing decision that shifts towards emptying earlier.

Table 17: Routing efficiency measured using km per stop under different planning horizons

Length of planning horizon (days)	Km/emptied container
1	0,43
2	0,38
3	0,36
4	0,35
5	0,35

Next to the performance on the KPIs, we also consider the differences in computation time. The computation time increases significantly as the length of the planning horizon increases as can be seen in Table 18.

Table 18: Computation time under different planning horizons (using an Intel i5 processor)

Length of planning horizon (days)	Average computation time (seconds)
1	34
2	183
3	467
4	906
5	1.478

Conclusion

Considering the performance of the different planning horizon lengths, using the efficiency frontier of Figure 20, a planning horizon of three days outperforms the other planning horizons. Moreover, using a planning horizon of three days has a routing efficiency close to the perceived best achievable, while still having a manageable computation time. For each of the configurations we choose to continue with a planning horizon of three days.

6.5 - Conclusion

This chapter shows the results of the numerical experiments conducted to evaluate the performance of the proposed solution approach and to study the effects of the experimental factors. The research question related to this chapter is: what is the expected outcome of the proposed planning methodology for the waste collection of Amsterdam?

Six experimental factors are considered: CFA method, penalty scaling factors, new cluster costs, AOP, implementing sensors, and the length of the planning horizon. An OFAT approach is adopted to find the configuration most suitable to the waste collection planning of Amsterdam. From the experiments, the following conclusions can be drawn on the effects of the experimental factors on the performance of the proposed solution approach:

- The differences in planning performance caused by the four CFA methods is minimal.
- The penalty scaling factors directly affect the timing decision by influencing the costs of changing the emptying day of containers and therefore has a large impact on the trade-off between the distance driven and the service level. Penalizing postponement more causes higher service levels, but also increases the distance driven per ton of waste collected, the opposite is true for penalizing early emptying more excessively.
- The costs for creating new clusters influences the routing efficiency and the decision on when to empty containers. Low costs for creating new clusters results in, on average, smaller cluster-dense containers which proves to improve the routing efficiency by reducing the distance driven per ton of waste collected. As a side-effect, the solution approach empties containers later, which results in a decrease in the service level. The opposite is true for high costs for creating new clusters.
- The AOP directly influences the timing decision by changing the desired emptying day and expected interval length between emptying the container. If a higher risk for overflowing is allowed, containers are emptied less, resulting in lower service levels, as well as less distance travelled per ton of waste.

- Installing sensors significantly improves the quality of solutions because of the increase of the quality of information used to make the planning. An average reduction of the distance travelled per ton of waste collected of 7,8% and a simultaneous increase of the service level by 3,9% can be achieved by installing sensors in the containers.
- The effect of the planning horizon is most notable on the routing efficiency. Longer planning horizons give the heuristic more flexibility, resulting in less distance travelled per collected container, however, as the number of containers emptied rises, this decrease is negated and the impact on the distance travelled per ton of waste is undone. Longer planning horizons also cause longer computation times. A planning horizon of three days is judged to be most suitable.

The previously discussed effects of the experimental parameters all influence the trade-off between the two KPIs: the distance driven per ton of waste and the service level offered to the inhabitants of Amsterdam. Different configurations of the experimental factors can be used to prioritize one of the two KPIs of the trade-off. Table 19 shows the extremes of prioritization of both KPIs and a more balanced trade-off. Moreover, the differences in planning performance to the current performance of Amsterdam is noted. In table 19, *Dist* denotes the distance travelled per ton of collected waste and *SL* denotes the achieved service level.

Table 19: Planning performance proposed solution approach with different priorities

Prioritization trade-off	Proposed solution approach		Δ to performance of Amsterdam	
	Dist	SL	Dist	SL
Prioritize distance	1,75	0,724	-18,2%	-10,0%
Balanced priorities	1,88	0,815	-12,1%	+1,4%
Prioritize service level	2,37	0,905	+10,7%	+12,6%

Chapter 7 – Conclusion and recommendations

This chapter summarizes the conclusions of this research and answers the main research question (Section 7.1). Moreover, several recommendations are listed for the municipality of Amsterdam (Section 7.2) and the discussion explores the applicability of this research to other problems and possibilities for further research (Section 7.3).

7.1 - Conclusion

This research started with the following core problem: the current static collection schedules and routes are unsuitable for the stochastic, dynamic demand for waste collection. Correspondingly, the objective of this research is to develop a dynamic planning methodology that is better suited to the characteristics of the waste collection in Amsterdam. This objective leads to the main research question: In which way and to what degree can the waste collection planning of Amsterdam be improved by using dynamic scheduling algorithms?

To start to answer this question, the current state of the waste collection system and its planning is studied. The analysis of the current situation shows that:

- Amsterdam's waste collection system consists of containers, wharfs, satellite facilities, and waste processing facilities, each of which has its own characteristics.
- The current schedules are static, fixed, and based on estimations of the required emptying frequency.
- The most relevant performance indicators for evaluating the planning performance are the influenceable collection costs, number of kilometers driven, and the service level.

Based on the exploration of the current situation and the problems faced in Amsterdam, a literature review is conducted to find relevant research papers. Several routing problems are discussed, of which the IRP is determined as the most suitable representation of the waste collection problem in Amsterdam. A lot of aspects from existing research are applicable to the problem of Amsterdam, but no paper covered all our requirements. Therefore, a new solution approach is proposed that consists of three phases that are executed in a rolling horizon framework:

1. Container selection, where containers that are expected to be relevant within the pre-determined planning horizon are selected.
2. Day assignment, where the selected containers are assigned to days of the planning horizon.
3. Route construction, where routes are constructed to collect all assigned containers.

A simulation model is used to perform numerical experiments to evaluate the performance of the proposed solution approach in comparison to the current planning methodology of Amsterdam. Moreover, the possibility of installing sensors into all containers, to mitigate the stochasticity inherent to the demand for waste collection, is studied. The results of the numerical experiments, as shown in Figure 21, show that the planning performance can be improved considerably by using a more dynamic planning methodology. The results are shown on an efficiency frontier as there is a trade-off between the distance travelled and offered service level. Decision makers should decide which point on the efficiency frontier best suits their objectives. The results show that a travel distance reduction of 12% is possible without installing sensors, and 21% with sensors, without reducing the service level offered to Amsterdam's inhabitants.

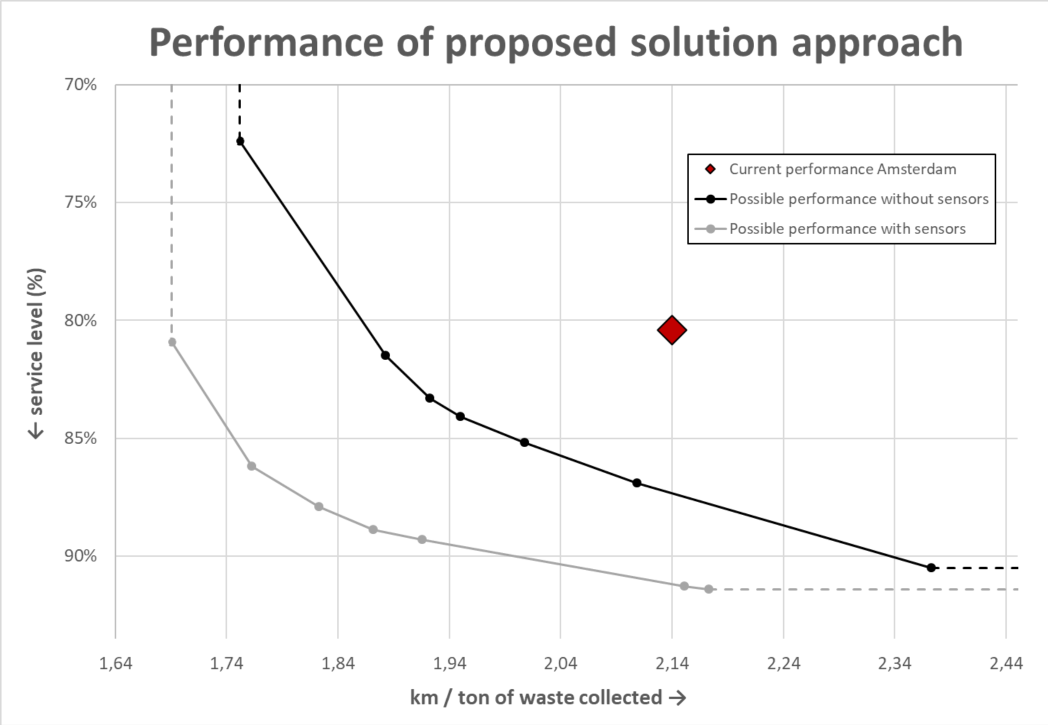

Figure 21: Possible performance of the proposed solution approach compared to Amsterdam

7.2 - Recommendations

Based on the conclusions of this research, we formulate several recommendations to the municipality of Amsterdam in this section. Moreover, we identify some areas in which the municipality can continue research to keep improving its waste collection planning.

Our first and foremost recommendation for the municipality of Amsterdam is to implement our proposed solution approach to plan its waste collection. More specifically, we recommend to use one of the configurations that lies on, or is at least close to, the efficiency frontier. The decision makers of Amsterdam should decide which of the efficient configurations best suits their objectives and use that configuration. The improvement of the proposed solution approach over the current planning methodology is not solely in the performance of the KPIs. The proposed solution approach also gives the possibility to consciously make decisions on planning aspects such as the acceptable risk of overflowing and the trade-off between the two conflicting KPIs: distance and service level.

Moreover, the municipality is recommended to consider the implementation of sensors into all or part of its containers. The potential savings, as discussed in Chapter 6, are significant, but implementing sensors in all containers would also involve considerable investment from the municipality. It is not necessary to implement sensors in all containers to achieve improvements. It is also possible to make a selection of containers, for example, those with high variances in deposits, which to equip with sensors.

If any investment into sensors is deemed too large, we recommend the municipality of Amsterdam to improve the quality of their data collection and analysis. With more data availability and more elaborate data analysis, it is expected that the fill level predictions of containers can improve considerably. This should include factors such as seasonality and expected growth of waste collection demand. The better the quality of the data, the better the quality of the fill level predictions and thus the planning performance.

Our last recommendation to the municipality is to initiate further research into an aspect that is largely omitted from this research on account of the limited time: dynamic decisions during the day. The current proposed solution considers incoming information, such as the overflowing of containers, only at the end of the day. This information can also be processed during the day, immediately upon receiving the information. Another option for dynamic decision making is to account for real-time traffic information to reconsider the shortest path to containers or even to reconsider the day assignment of containers. The implementation and potential benefit of these measures deserve further research.

To summarize, our recommendations to the municipality of Amsterdam are to:

- Implement the proposed solution approach using one of the preferred configurations,
- Implement sensors into the containers,
- Improve the quality of data collection to enable better fill level predictions,
- Initiate further research into the benefits of additional planning dynamicity.

7.3 - Discussion

During the discussion, we briefly discuss the applicability of the proposed solution approach to problems outside the environment its currently tested in (Section 7.3.1), the limitations of this research (Section 7.3.2), and give some recommendations on further research (Section 7.3.3).

7.3.1 - Wider applicability of solution approach

Currently, the proposed solution approach is solely implemented and tested in Amsterdam’s Zuidoost district. However, it is expected that its implementation to the entirety of Amsterdam is relatively trivial. The exact same data processing steps, for example, to find the parameters of the distributions of each container’s waste deposits, can be followed. The only additional difficulty is to extend the solution approach to decide from which wharf each route should start.

Furthermore, the applicability of the proposed solution approach is not limited to the waste collection planning of Amsterdam. We expect that similar improvements are possible in other municipalities that are currently utilizing the same planning techniques as used in Amsterdam. Moreover, other IRP problems, not related to waste collection, may also benefit from applying our solution approach. The difference of the solution approach proposed in this thesis to existing IRP solution methods is the integrated way in which decisions regarding the time and space dimensions are made. However, further research should first be conducted before the benefit for other types of IRPs can be confirmed definitively.

However, the generalizability of the solution approach is promising as it allows for calibration of various aspects in such a way that it can adjust to the characteristics of the problem to which it is applied. For example, it is expected that if the geographical dispersion of customers changes, the costs for creating new clusters should be adjusted accordingly. It is believed that the calibration approach such as followed in Chapter 6 can be applied for each new application. Because of this flexibility of the proposed solution approach, it is believed that it can be applied in a variety of settings.

7.3.2 - Limitations of research

The foremost limitation of this research is the difficulty with which the performance of the proposed solution approach can be compared with that of the current planning of Amsterdam. This limitation is self-imposed, as one of the simplifications of the simulation model is to allow household waste vehicles to utilize both the satellite facility and the waste processor in the Zuidoost district. This made the problem more interesting and relevant, but also complicated the comparison with the actual situation, as this is not the case in reality. Moreover, the lack of available data on the current service level of Amsterdam further complicated comparison with reality, as this is one of the two most important KPIs. To solve this problem, the current planning methodology of Amsterdam is mimicked in the simulation model. However, this is only an approximation and does not fully capture all details of the planning and scheduling done in Amsterdam. Therefore, this research is not fully able to show the precise improvements made possible by the proposed solution approach.

Moreover, the focus of this research is on a novel way of simultaneously considering both the time and space dimensions of the IRP. Therefore, the implemented routing and improvement heuristic remain relatively elementary. Implementing more elaborate routing and improvement heuristics is expected to further improve the solution quality.

Another limitation is the limited contact with the municipality of Amsterdam during this research. This left some questions about the current planning methodology and prevented additional validation of the simulation model by subject experts, further complicating the comparison of the results of the experiments with the current planning performance.

7.3.3 - Further research

This thesis introduces a novel way of solving the IRP, while simultaneously considering the time and space dimensions of the IRP. This approach is based on approximations of the travel costs and costs of the timing of emptying a container, respectively represented by the cluster fitness approximation and timing penalty costs in our research. As these are merely rough approximations, it is an interesting topic of further research. Other approximation methods can be developed and researched, potentially better representing the costs of the associated decisions, improving the solution quality.

Moreover, as already discussed in Section 7.3.1, the applicability of the proposed solution approach to a wider range of IRP applications should be researched to confirm that this approach is suitable as a standard solution approach for the IRP, or if it is only applicable to waste collection problems.

List of abbreviations

An alphabetical list of all abbreviations used in this thesis is given below.

ALNS	Adaptive Large Neighborhood Search
C&W	Clark & Wright
CARP	Capacitated Arc Routing Problem
CIRP	Cyclic Inventory Routing Problem
CLRIP	Combined Location Routing and Inventory Problem
CVRP	Capacitated Vehicle Routing Problem
DES	Discrete-Event Simulation
GIS	Geographical Information System
Goudappel	Goudappel Group, a collaboration of companies, most notably Goudappel Coffeng and DAT.Mobility.
IP	Inventory Problem
IRP	Inventory Routing Problem
IRPSF	Inventory Routing Problem with Satellite Facilities
IRPT	Inventory Routing Problem with Transshipment
IRP-CM	Inventory Routing Problem with Continuous Moves
KPI	Key Performance Indicator
MCDM	Multiple Criteria Decision Making
MILP	Mixed Integer Linear Programming
MMIRP	Multi-Product Multi-Vehicle Inventory Routing Problem
RVND	Randomized Variable Neighborhood Descent
SA	Simulated Annealing
TOP	Team Orienteering Problem
TSP	Traveling Salesman Problem
VMI	Vendor Managed Inventory
VRP	Vehicle Routing Problem
VRPSF	Vehicle Routing Problem with Satellite Facilities
WCVRP	Waste Collection Vehicle Routing Problem
WKT	Well-Known Text

References

Aguilar-Savén, R. (2004). Business process modelling: Review and framework. *International Journal of Production Economics*, 129-149.

Aksen, D., Kaya, O., Salman, F., & Tüncel. (2014). An adaptive large neighborhood search algorithm for a selective and periodic inventory routing problem. *European Journal of Operational Research*, 413-426.

Andersson, H., Hoff, A., Christiansen, M., Hasle, G., & Løkketangen, A. (2010). Industrial aspects and literature survey: Combined inventory management and routing. *Computers & Operations Research*, 1515-1536.

Anily, S., & Bramel, J. (2004). An asymptotic 98.5%-effective lower bound on fixed partition policies for the inventory-routing problem. *Discrete Applied Mathematics*, 22-39.

Archetti, C., Bertazzi, L., Laporte, G., & Speranza, M. (2007). A Branch-and-Cut Algorithm for a Vendor-Managed Inventory-Routing Problem. *Transportation Science*, 382-391.

Baita, F., Ukovich, W., Pesenti, R., & Favaretto, D. (1998). Dynamic routing-and-inventory problems: a review. *Transportation Research Part A: Policy and Practice*, 585-598.

Baldacci, R., & Maniezzo, V. (2004). *Exact methods based on node routing formulations for arc routing problems.* University of Bologna: Department of Computer Science.

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

Beliën, J., de Boeck, L., & van Ackere, J. (2014). Municipal Solid Waste Collection and Management Problems: A Literature Review. *Transportation Science*, 78-102.

Bell, W., Dalberto, L., Fisher, M., Greenfield, A., Jaikumar, R., Kedia, P., . . . Prutzman, P. (1983). Improving the distribution of industrial gases with an on-line computerized routing and scheduling optimizer. *Interfaces*, 4-23.

Bellman, R. (1958). On a Routing Problem. *Quarterly of Applied Mathematics*, 87-90.

Benjamin, A., & Beasley, J. (2010). Metaheuristics for the waste collection vehicle routing problem with time windows, driver rest period and multiple disposal facilities. *Computers & Operations Research*, 2270-2280.

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

Bertazzi, L., Savelsbergh, M., & Speranza, M. (2008). Inventory routing. In B. Golden, S. Raghavan, & E. Wasil, *The vehicle routing problem: latest advances and new challenges* (pp. 49-72). Boston: Springer.

Blum, C., & Roli, A. (2003). Metaheuristics in combinatorial optimization: Overview and conceptual comparison. *ACM computing surveys (CSUR)*, 268-308.

Braekers, K., Ramaekers, K., & van Nieuwenhuyse, I. (2016). The vehicle routing problem: State of the art classification and review. *Computers & Industrial Engineering*, 300-313.

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

Campbell, A., Clarke, L., Kleywegt, A., & Savelsbergh, M. (1998). The inventory routing problem. In *Fleet management and logistics* (pp. 95-113). Boston: Springer.

Cheng, C., Qi, M., Wang, X., & Zhang, Y. (2016). Multi-period inventory routing problem under carbon emission regulations. *International Journal of Production Economics*, 263-275.

Cheng, C., Yang, P., Qi, M., & Rousseau, L. (2017). Modeling a green inventory routing problem with a heterogeneous fleet. *Transportation Research*, 97-112.

Chiou, S. (2005). Integrating the inventory management and vehicle routing problems for congested urban logistics network. *Journal of the Eastern Asia Society for Transportation Studies*, 3038-3051.

Chitsaz, M., Divsalar, A., & Vansteenwegen, P. (2016). A two-phase algorithm for the cyclic inventory routing problem. *European Journal of Operations Research*, 410-426.

Coelho, L., & Laporte, G. (2013). A branch-and-cut algorithm for the multi-product multi-vehicle inventory-routing problem. *International Journal of Production Research*, 7156-7169.

Coelho, L., Cordeau, J., & Laporte, G. (2012). The inventory-routing problem with transshipment. *Computers & Operations Research*, 2537-2548.

Coelho, L., Cordeau, J., & Laporte, G. (2014). Thirty Years of Inventory Routing. *Transportation Science*, 1-19.

Cordeau, J., Gendreau, M., Laporte, G., Potvin, J., & Semet, F. (2002). A guide to vehicle routing heuristics. *Journal of Operational Research Society*, 512-522.

Cordeau, J., Laganà, D., Musmanno, R., & Vocaturo, F. (2015). A decomposition-based heuristic for the multiple-product inventory-routing problem. *Computers & Operations Research*, 153-166.

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

Das, S., & Bhattacharyya, B. (2015). Optimization of municipal solid waste collection and transportation routes. *Waste Management*, 9-18.

Delling, D., Sanders, P., Schultes, D., & Wagner, D. (2009). Engineering Route Planning Algorithms. In J. Lerner, D. Wagner, & K. Zweig, *Algorithmics of Large and Complex Networks* (pp. 117-139). Berlin: Springer.

Di Somma, M. (2016). *Optimal operation planning of distributed energy systems through multi-objective approach: A new sustainability-oriented pathway.*

Dijkstra, E. (1959). A note on two problems in connection with graphs. *Numerische Mathematik 1*, 269-271.

Dror, M., Ball, M., & Golden, B. (1985). A computational comparison of algorithms for the inventory routing problem. *Annals of Operations Research*, 1-23.

Federgruen, A., & Zipkin, P. (1984). A combined vehicle routing and inventory allocation problem. *Operations Research*, 1019-1037.

Ferreira, J., Quintas, A., Oliveira, J., Pereira, G., & Dias, L. (2014). Solving the team orienteering problem: developing a solution tool using a genetic algorithm approach. In *Soft Computing in Industrial Applications* (pp. 365-375). Cham: Springer.

Geetha, S., Poonthalir, G., & Vanathi, P. (2009). Improved k-means algorithm for capacitated clustering problem. *INFOCOMP*, 52-59.

Goldberg, A., Kaplan, H., & Werneck, R. (2006). Reach for A*: Efficient point-to-point shortest path algorithms. *Proceedings of the Eighth Workshop on Algorithm Engineering and Experiments (ALENEX)* (pp. 129-143). Society for Industrial and Applied Mathematics.

Han, H., & Ponce-Cueto, E. (2015). Waste collection vehicle routing problem: Literature review. *PROMET - Traffic & Transportation*, 345-358.

Heerkens, H., & van Winden, A. (2012). *Geen Probleem.* Buren: Business School Nederland.

Hertz, A., Laporte, G., & Mittaz, M. (2000). A tabu search heuristic for the capacitated arc routing problem. *Operations Research*, 129-135.

Huang, B., Fery, P., Xue, L., & Wang, Y. (2008). Seeking the Pareto front for multiobjective spatial optimization problems. *International Journal of Geographical Information Science*, 507-526.

Huang, B., Wu, Q., & Zhan, F. (2007). A shortest path algorithm with novel heuristics for dynamic transportation networks. *International Journal of Geographical Information Science*, 625-644.

Jaillet, P., Bard, J., Huang, L., & Dror, M. (2002). Delivery cost approximations for inventory routing problems in a rolling horizon framework. *Transportation Science*, 292-300.

Jewel. (2017). *Jewel bespaart Utrecht veel moeite*. Retrieved August 07, 2018, from https://www.jewel.eu/en/portfolio_item/jewel-bespaart-utrecht-tijd-en-moeite/

Johansson, O. (2006). The effect of dynamic scheduling and routing in a solid waste management system. *Waste Management*, 875-885.

Kim, B., Kim, S., & Sahoo, S. (2006). Waste collection vehicle routing problem with time windows. *Computers & Operations Research*, 3624-3642.

Kleywegt, A., Nori, V., & Savelsbergh, M. (2002). The Stochastic Inventory Routing Problem with Direct Deliveries. *Transportation Science*, 94-118.

Kumar, S., & Panneerselvam, R. (2012). A Survey on the Vehicle Routing Problem and Its Variants. *Intelligent Information Management*, 66-74.

Laporte, G. (1992). The Vehicle Routing Problem: An overview of exact and approximate algorithms. *European Journal of Operational Research*, 345-358.

Law, A. (2015). *Simulation Modeling and Analysis.* New York: McGraw-Hill Education.

Lenstra, J., & Kan, A. (1975). Some simple applications of the travelling salesman problem. *Journal of the Operational Research Society*, 717-733.

Levinson, D., & El-Geneidy, A. (2009). The minimum circuity frontier and the journey to work. *Regional Science and Urban Economics*, 732-738.

Li, J., Borenstein, D., & Mirchandani, P. (2007). A decision support system for the single-depot vehicle rescheduling problem. *Computers & Operations Research*, 1008-1032.

Liu, S., & Lin, C. (2005). A heuristic method for the combined location routing and inventory problem. *The International Journal of Advanced Manufacturing Technology*, 372-381.

Medvedev, A., Fedchenkov, P., Zaslavsky, A., Anagnostopoulos, T., & Khoruzhnikov, S. (2015). Waste management as an IoT-enabled service in smart cities. In *Internet of Things, Smart Spaces, and Next Generation Networks and Systems* (pp. 104-115). Springer.

Mes, M. (2012). Using simulation to assess the opportunities of dynamic waste colelction. In *Use cases of discrete event simulation* (pp. 277-307). Berlin, Heidelberg: Springer.

Mes, M., Schutten, M., & Rivera, A. (2014). Inventory routing for dynamic waste collection. *Waste Management*, 1564-1576.

Moin, N., & Salhi, S. (2007). Inventory routing problems: a logistical overview. *Journal of the Operational Research Society*, 1185-1194.

Municipality of Amsterdam. (2015). *Afvalketen in beeld.*

Municipality of Amsterdam. (2016). *Uitvoeringsplan Afval: Grondstoffen uit Amsterdam.*

Municipality of Amsterdam. (2018a). *Gunningsleidraad met uitnodiging tot Eerste Inschrijving.*

Municipality of Amsterdam. (2018b). *Kerncijfers Amsterdam 2018.* Amsterdam: Onderzoek, Informatie en Statistiek.

Municipality of Amsterdam. (2018c). *Selectieleidraad van de Europese aanbesteding voor in Innovatiepartnerschap ten behoeve van optimalisatie van afvalinzamelingsprocessen in de gemeente Amsterdam.* Amsterdam.

Nilsson, C. (2003). *Heuristics for the traveling salesman problem.* Linkoping University.

Nuortio, T., Kytöjoki, J., Niska, H., & Bräysy, O. (2006). Improved route planning and scheduling of waste collection and transport. *Expert Systems with Applications*, 223-232.

Oppen, J., & Løkketangen, A. (2006). Arc routing in a node routing environment. *Computers & Operations Research*, 1033-1055.

Papadimitriou, C. (2003). *Computational complexity (pp. 260-265).* John Wiley and Sons Ltd.

Pearn, W., Assad, A., & Golden, B. (1987). Transforming arc routing into node routing problems. *Computers & Operations Research*, 285-288.

Peres, I., Repolho, H., Martinelli, R., & Monteiro, N. (2017). Optimization in inventory-routing problem with planned transshipment: A case study in the retail industry. *International Journal of Production Economics*, 748-756.

Qu, W., Bookbinder, J., & Iyogun, P. (1999). An integrated inventory-transportation system with modified periodic policy for multiple products. *European Journal of Operational Research*, 254-269.

Ramkumar, N., Subramanian, P., Narendran, T., & Ganesh, K. (2012). Mixed integer linear programming model for multi-commodity multi-depot inventory routing problem. *Operations Research & Decision Theory*, 413-429.

Robusté, F., Daganzo, C., & Souleyrette, R. (1990). Implementing vehicle routing models. *Transportation Research*, 263-286.

Samanlioglu, F. (2013). A multi-objective mathematical model for the industrial hazardous waste location-routing problem. *European Journal of Operational Research*, 332-340.

Sari, K. (2007). Exploring the benefits of vendor managed inventory. *International Journal of Physical Distribution & Logistics Management*, 529-545.

Savelsbergh, M., & Song, J. (2008). An optimization algorithm for the inventory routing problem with continuous moves. *Computers & Operations Research*, 2266-2282.

Stewart, T. (1992). A critical survey on the status of multiple criteria decision making theory and practice. *Omega*, 569-586.

Ticha, H., Absi, N., Feillet, D., & Quilliot, A. (2017). *A solution method for the multi-destination, bi-objective shortest path problem.* Ecole des Mines de Saint Etienne.

Tommiska, M., & Skytta, J. (2001). Dijkstra's shortest path routing algorithm in reconfigurable hardware. *International Conference on Field Programmable Logic and Applications* (pp. 653-657). Berlin: Springer.

Vicentini, F., Giusti, A., Rovetta, A., Fan, X., He, Q., Zhu, M., & Liu, B. (2009). Sensorized waste collection container for content estimation and collection optimization. *Waste Management*, 1468-1472.

Wang, X., & Kopfer, H. (2015). Rolling horizon planning for a dynamic collaborative routing problem with full-truckload pickup and delivery requests. *Flexible Services and Manufacturing Journal*, 509-533.

Webster, J., & Watson, R. (2002). Analyzing the past to prepare for the future: Writing a literature review. *Management Information Systems Quarterly*, xiii-xxiii.

Woeginger, G. (2003). Exact algorithms for NP-hard problems: A survey. In *Combinatorial Optimization - Eureka, you shrink!* (pp. 185-207). Berlin: Springer.

Xue, W., & Cao, K. (2016). Optimal routing for waste collection: a case study in Singapore. *International Journal of Geographical Information Science*, 554-572.

Appendix 1 - Problem cluster

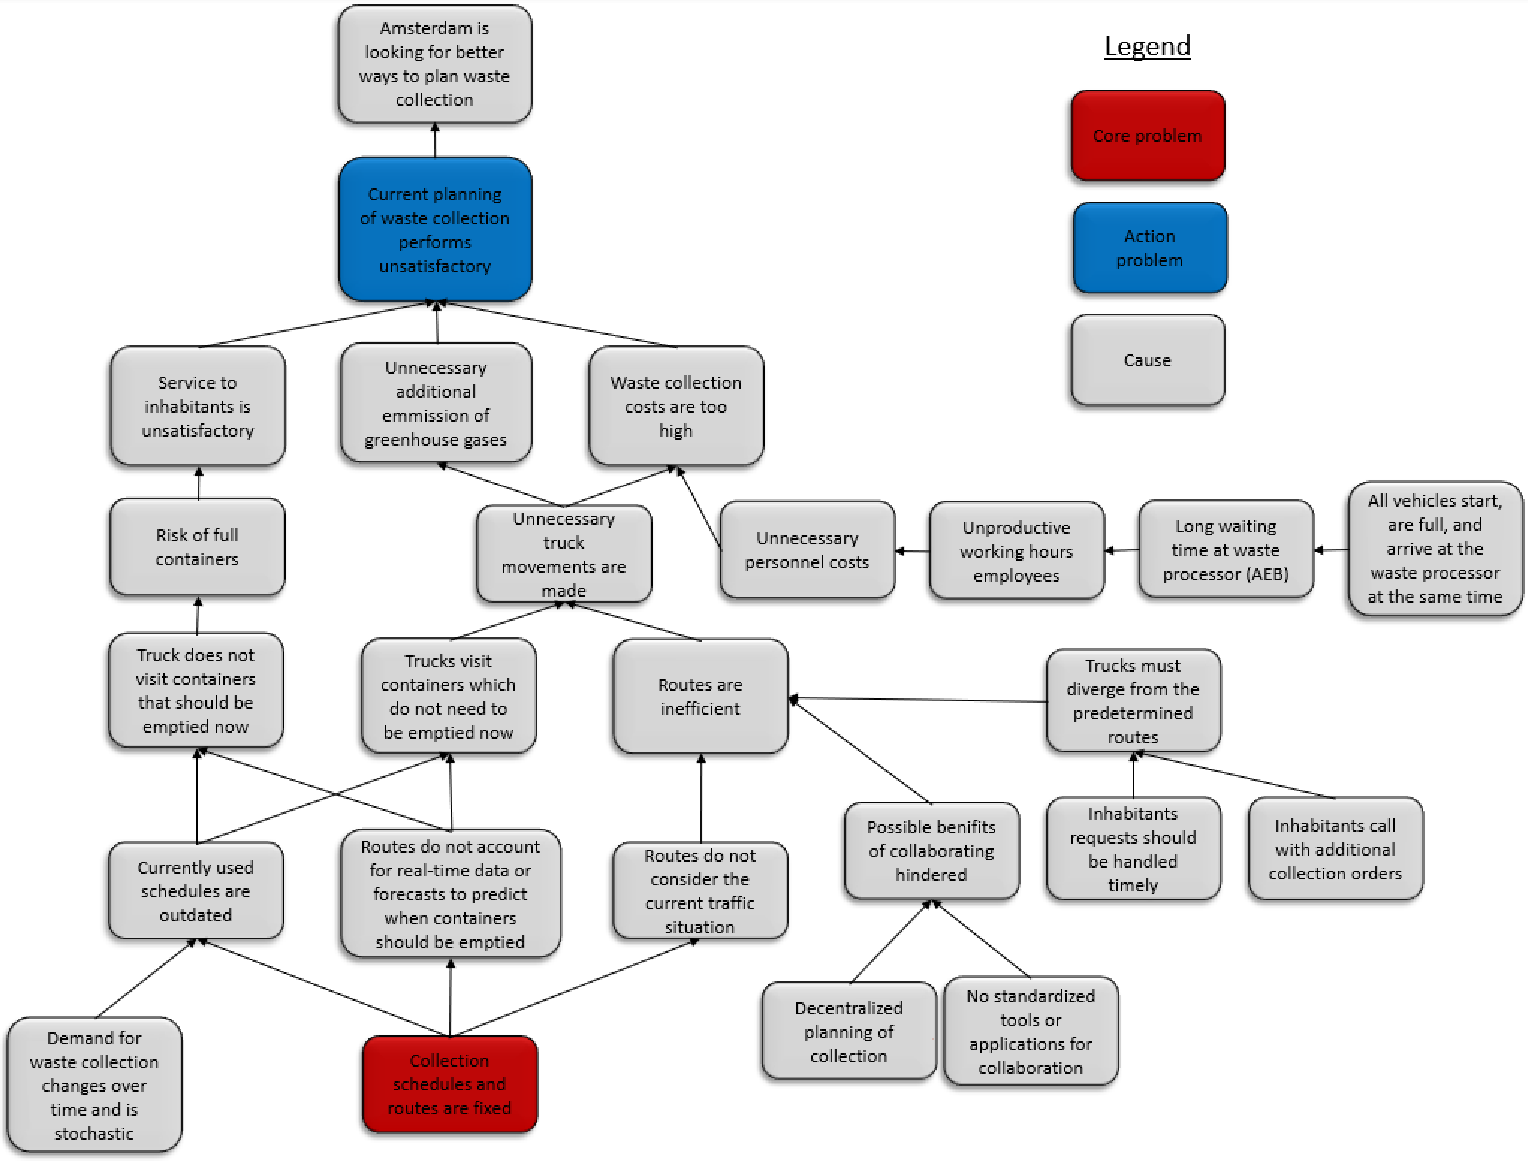

Appendix 2 - Distribution of generated waste per waste fraction

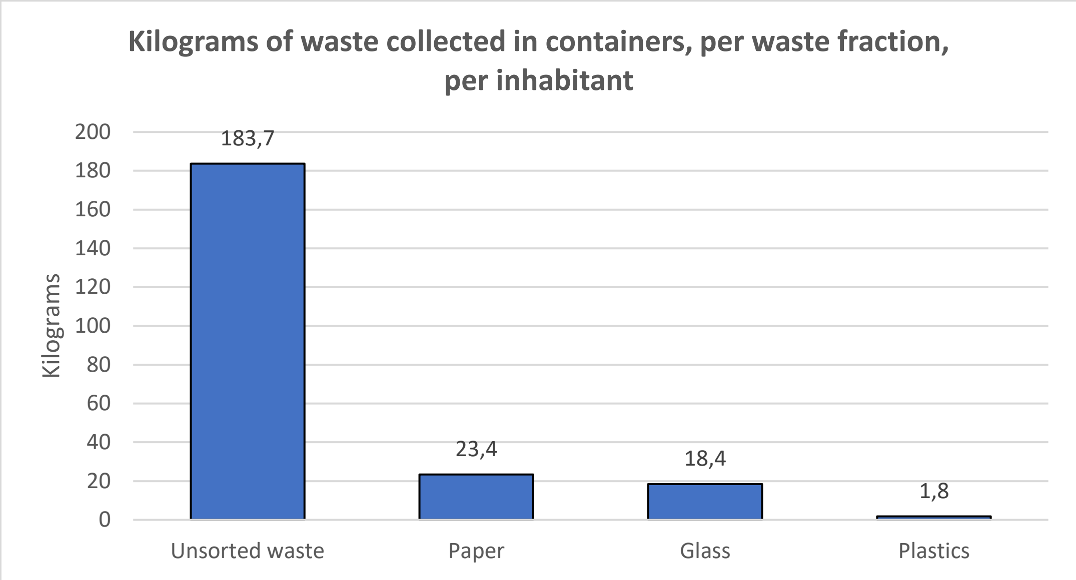

Figure 5: Kilograms of waste collected in containers, per waste fraction, per inhabitant, per year

Appendix 3 - Waste collection related complaints in Amsterdam

Figure 6: Complaints related to waste collection (in 2014)

Appendix 4 - Classification schemes from current literature reviews

Authors (year)	Classification on the topics of…
Baita, Ukovich, Pesenti, & Favaretto (1998)	Topology, number of items, type of demand, decision domain, constraints, costs, proposed solution approaches
Moin & Salhi (2007)	Single period, multiperiod, infinite horizon models, and stochastic demand patterns
Andersson, Hoff, Christiansen, Hasle, & Løkketangen (2010)	Time, demand, topology, routing, inventory, fleet composition, fleet size
Bertazzi & Speranza (2012)	Shipping times, planning horizon, structure of distribution policy, objective of policy, and decision space
Coelho, Cordeau, & Laporte (2014)	Time horizon, structure, routing, inventory policy, inventory decisions, fleet composition, and fleet size. (referencing Andersson et al. (2010))

Appendix 5 - Supplementary algorithms

Algorithm 2.1 - Cluster overflowed containers algorithm

Algorithm that is used during the day assignment phase where the overflowed containers of yesterday are scheduled first.

Algorithm 2.1: Cluster overflowed containers algorithm

Algorithm 2.1: Cluster overflowed containers algorithm

Input: set of overflowed containers yesterday (OC)
Output: set of clusters (CL)
Parameters: costs of adding container i to cluster j (ac_{ij}), costs of creating a new cluster for container i (nc_i)

```
0.  Initialize, CL = ∅
1.  While OC ≠ ∅ do
2.  |   For each container i in set OC
3.  |   |   For all clusters j in set CL
4.  |   |   |   calculate ac_ij
5.  |   |   |   remember cheapest ac_ij
6.  |   End for
7.  |   If cheapest ac_ij < nc_i then
8.  |   |   add container i to cluster j
9.  |   Else
10. |   |   create new cluster with container i
11. |   End if
12. End while
13. Return CL
```

K-means algorithm

Algorithm used to pre-process a list of container to make suitable clusters for route creation.

Algorithm: K-means algorithm

Input: set of containers (C), set of day-assignment clusters (DA-CL)
Output: set of pre-processed clusters (PP-CL)
Parameters: capacity vehicles, buffer vehicles, number of clusters (k), container-to-cluster priority

```
0.  k = total load / (vehicle capacity * (1 – vehicle buffer))
1.  Find k starting points for clusters based on largest clusters in DA-CL
2.  For 10 iterations do
3.  |   Calculate priority for each container to closest cluster using Equation 8
4.  |   While unassigned containers ≠ 0 do
5.  |   |   Find container with highest priority
6.  |   |   If possible to assign regarding cluster capacity then
7.  |   |   |   Assign container to closest cluster
8.  |   |   Else if not possible to assign container with highest priority then
9.  |   |   |   Recalculate container priority with second closest cluster
10. |   |   |   Return to step 4.
11. |   |   End if
12. |   End while
13. |   Recalculate coordinates cluster centroids
14. End for
15. Return PP-CL
```

$$priority_i = \frac{expected\ fill\ level\ container}{distance\ to\ closest\ cluster} \tag{9}$$

2-opt algorithm

Improvement algorithm that is used to improve the initially created routes. Note that route[i] means the i-th place in the routing sequence.

Algorithm: 2-opt algorithm

Input: set of routes (R)
Output: set of routes (R)
Parameters: set of tours (T), convergence, iterators i and j

```
For each route in R do
|   Create list of tours T by splitting routes into tours separated by disposal facilities
|   For each tour in T do
|   |   Convergence = false
|   |   While convergence = false do
|   |   |   Convergence = true
|   |   |   For i = begin tour to i <= end tour – 2 do
|   |   |   |   For j = i + 1 to j <= end tour – 1 do
|   |   |   |   |   Add route[begin tour] to route[i-1] to sequence
|   |   |   |   |   Add route[i] to route[j] to sequence in reverse order
|   |   |   |   |   Add route[j+1] to route[end tour] to sequence
|   |   |   |   |   If improvement to tour then
|   |   |   |   |   |   Keep changes
|   |   |   |   |   |   Convergence = false
|   |   |   |   |   If no improvement then
|   |   |   |   |   |   Reverse changes
|   |   |   |   |   End if
|   |   |   |   End for
|   |   |   End for
|   |   End while
|   End for
End for
```

Appendix 6 - Assumptions and simplifications simulation model

This appendix lists the assumptions and simplifications made during the modeling of the simulation model of Amsterdam Zuidoost. In addition, the choice to change the functionality of the waste processing facility of Zuidoost is justified.

Assumptions

- Emptying time container = 3 minutes
- Processing time disposal facilities = 15 minutes
- Weight in kg per cubic meter of household waste = 100 kg
- Vehicle capacity = 9.000 kg
- Average vehicle speed = 15 km/h
- Distances between locations are Euclidean distances * 1.2 (Levinson and El-Geneidy, 2009)
- Container fill rates are modeled using a Gamma-distribution
- Container fill rates do not change when a container is full
- Unlimited extra deposits can be made to a container, even if full
- Containers can only be emptied completely, partial emptying is impossible
- A sufficient number of collection vehicles and operating personnel is always available
- Collection vehicles can spend 6.25 hours on work related activities

Simplifications

- Only considering household waste
- All vehicles are homogeneous
- Waste is only deposited by inhabitants at the end of the day
- Waste is only collected by the municipality at the start of the day
- A vehicle buffer of 25% is implemented and resulting routes are assumed to always be feasible
- Travel times between locations are deterministic and determined by the distance and the average speed
- No lunch break is scheduled, even though the required time is subtracted from the total working time

Changes of functionality waste processing facility

In reality, the waste processing facility located in Amsterdam Zuidoost is dedicated to processing paper and is thus not compatible with the household waste fraction. Therefore, in reality, all household waste of Zuidoost is transported to the satellite facility, from where it is transported to the north of Amsterdam, where the household waste processor is located. However, we want to solve the IRPSF, as this is the problem faced by the Amsterdam. The IRPSF has the additional decision between disposing waste at a satellite facility or a waste processing facility, this choice would be trivial if no changes are made. Therefore, in the simulation model, we pretend that the waste processing facility is compatible with the household waste fraction. To make the decision between both disposal facilities relevant, we introduce a fee associated with using the satellite facility. This is realistic as the waste disposed at the satellite facility should still be transported, albeit by different, cheaper methods, to the waste processing facilities.

The fee incurred for each container visiting the satellite facility is a penalty of 2.6 kilometers, which is based on the distance between the satellite facility and waste processing facility times a correction factor of 1.2 prescribed for urban areas by Levinson and El-Geneidy (2009). This fee keeps the decision between the two disposal facilities relevant.

Appendix 7 - Replication deletion approach calculations

Warmup period

Five replications are performed, each of sufficient length, in this case 50 days. Using windows of 5 and 10 days, the Welch's graphical procedure is applied as seen in the following figure. The chosen length of the warmup period is 25 days.

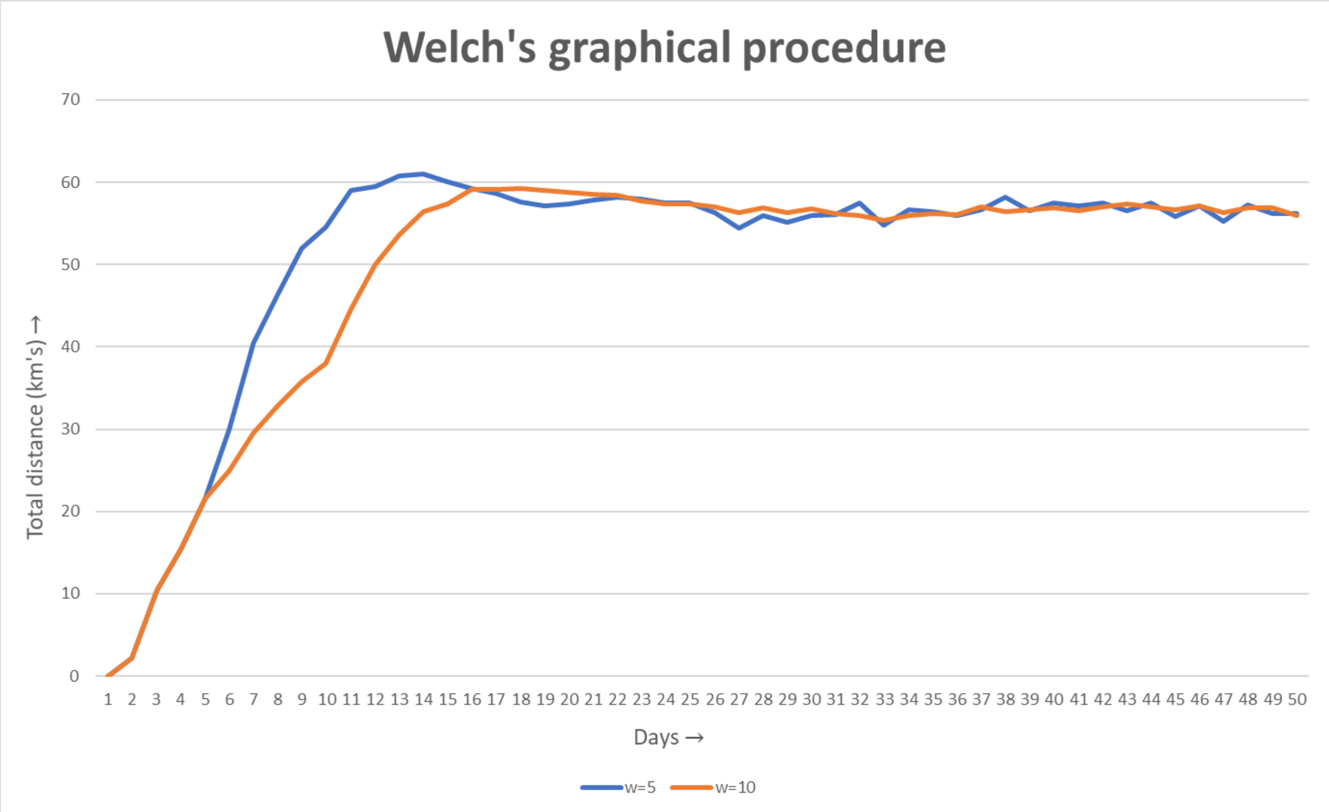

Number of replications

Using the sequential procedure of Law (2015) with an alpha and relative error of 0,05, the results are as follows, determining the required number of replications to be 3.

noRepl	TotDist	Average	StDev	Tstatistic	Delta	Error	Validation		alpha	0,05
1	5677,90	5677,90	0	0	0	0	false		relative error	0,05
2	5543,55	5610,73	95,00	12,7062	853,5348	0,152125	false			
3	5585,48	5602,31	68,74	4,302653	170,7548	0,030479	true			
4	5624,83	5607,94	57,24	3,182446	91,08587	0,016242	true			
5	5643,29	5615,01	52,03	2,776445	64,60667	0,011506	true			

Appendix 8 - Implementation of the Amsterdam's current planning methodology

This appendix discusses the implementation of Amsterdam's current planning methodology in the simulation model. Amsterdam's current planning methodology is based on fixed emptying frequencies. This means that the interval between emptying a container two subsequent times is constant. To mimic the current planning of Amsterdam, we should therefore know the emptying frequencies of each container. As this information is not available, we try to approximate them for each container. This is done by relating the emptying frequency to the average fill level upon emptying of a container. Equation 11 shows that the two variables are directly related, as the other parameters in the equation are constants. Therefore, if the average fill rate achieved in reality can be approached in the simulation model, the average emptying frequency is also approximated.

$$Average\ emptying\ frequency/week = \frac{average\ waste\ deposited/container/week}{average\ fill\ level\ containers * \#containers} \quad (11)$$

This is done using the technique also used in Section 4.3 to find the interval length between emptying, using the container fill speed characteristics and the AOP. As the container fill speed characteristics are given, the AOP is altered until the average fill level of containers upon emptying under the current planning methodology is approximated by that of the simulation model. The current the average fill level of containers upon emptying in Amsterdam Zuidoost is 53%. Experiments show that using an AOP of $0{,}15$ the average fill rate of containers upon emptying is 50%.

Using this AOP, the constant interval between emptying a container twice is determined which is then used to schedule containers. The timing decision for each container is solely dependent on the emptying interval and the same routing heuristics are used to solve the VRPSF, meaning that all improvements upon the current performance are made during the day assignment phase of the proposed solution approach.

Mimicking the planning methodology of Amsterdam in the implemented simulation model, while ensuring similar emptying frequencies for each container, we achieve the following results: $2{,}14$ kilometers driven per collected ton of waste and a service level of $80{,}4\%$. The municipality currently estimates that, for the collection of household waste in Amsterdam Zuidoost, it drives an average of $2{,}1$ kilometers per ton of waste, so the outcome of the simulation are expected to be realistic.

Appendix 9 - Cost function timing penalty costs

This appendix shows an example of the timing penalty costs function for a random container with the following characteristics:

- Shape (week) = 1.6
- Scale (week) = 316.12
- AOP = 0.25
- SCRC = 6.9
- CFA = 0.45
- DED = 5

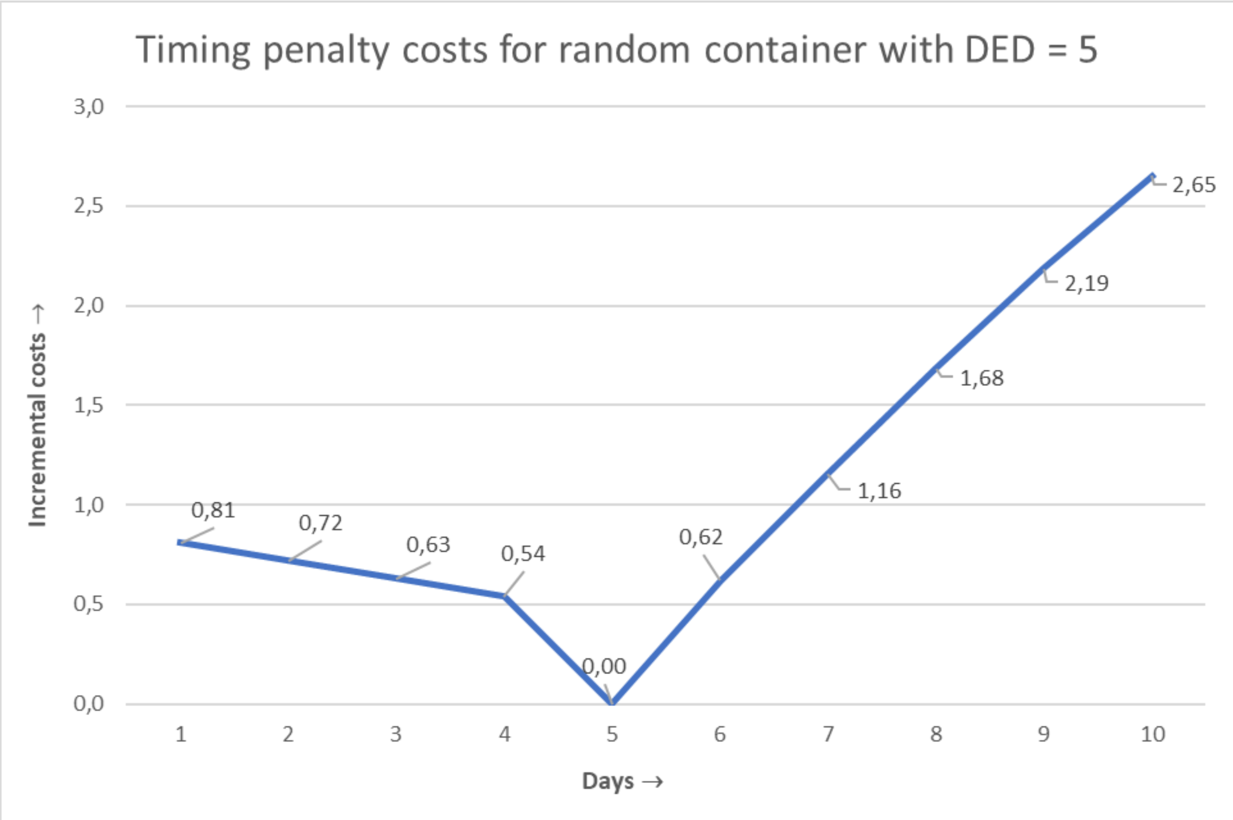

Appendix 10 - Full results of sensors experiment

CFA method	Penalty scaling	AOP	Sensor information: Without sensors Dist	SL	With sensors Dist	SL	Δ (%) Dist	SL
CC	Medium	0,05	2,26	0,886	2,26	0,886	-9,8%	+1,4%
		0,10	2,16	0,868	1,92	0,889	-11,3%	+2,4%
		0,15	2,07	0,858	1,82	0,879	-11,8%	+2,5%
		0,20	1,95	0,841	1,79	0,861	-8,2%	+2,4%
		0,30	1,89	0,796	1,73	0,835	-8,3%	+4,9%
CP	Low, low	0,05	2,11	0,869	2,02	0,885	-4,2%	+1,9%
		0,10	2,00	0,849	1,90	0,877	-4,7%	+3,3%
		0,15	1,92	0,833	1,82	0,870	-5,6%	+4,4%
		0,20	1,87	0,800	1,78	0,851	-4,8%	+6,4%
		0,30	1,81	0,757	1,72	0,824	-4,8%	+8,9%
	Medium	0,05	2,37	0,905	2,15	0,911	-9,4%	+0,6%
		0,10	2,24	0,884	2,01	0,899	-10,6%	+1,7%
		0,15	2,15	0,868	1,92	0,893	-11,0%	+2,9%
		0,20	2,01	0,852	1,83	0,878	-8,7%	+3,0%
		0,30	1,90	0,803	1,76	0,848	-7,7%	+5,6%
CI	Low, low	0,05	2,09	0,862	2,02	0,881	-3,4%	+2,2%
		0,10	1,97	0,839	1,89	0,868	-4,2%	+3,4%
		0,15	1,88	0,815	1,81	0,856	-3,9%	+5,1%
		0,20	1,85	0,780	1,74	0,840	-5,7%	+7,7%
		0,30	1,75	0,724	1,69	0,809	-3,5%	+11,7%
	Medium	0,05	2,36	0,900	2,17	0,914	-8,0%	+1,6%
		0,10	2,18	0,873	1,99	0,900	-8,6%	+3,0%
		0,15	2,06	0,855	1,87	0,889	-9,3%	+3,9%
		0,20	1,92	0,827	1,80	0,872	-6,0%	+5,5%
		0,30	1,83	0,758	1,73	0,838	-5,6%	+10,5%
	High, high	0,05	2,37	0,904	2,15	0,913	-9,4%	+0,9%
		0,10	2,28	0,885	2,01	0,898	-11,9%	+1,5%
		0,15	2,19	0,870	1,92	0,889	-12,3%	+2,3%
		0,20	2,09	0,859	1,83	0,880	-12,1%	+2,4%
		0,30	1,95	0,820	1,77	0,847	-9,3%	+3,2%