

Rich Vehicle and Inventory Routing Problems with Stochastic Demands

THÈSE N° 8009 (2017)

PRÉSENTÉE LE 24 NOVEMBRE 2017

À LA FACULTÉ DE L'ENVIRONNEMENT NATUREL, ARCHITECTURAL ET CONSTRUIT

LABORATOIRE TRANSPORT ET MOBILITÉ

PROGRAMME DOCTORAL EN GÉNIE CIVIL ET ENVIRONNEMENT

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

POUR L'OBTENTION DU GRADE DE DOCTEUR ÈS SCIENCES

PAR

Iliya Dimitrov MARKOV

acceptée sur proposition du jury:

Prof. C. J. D. Fivet, président du jury

Prof. M. Bierlaire, Prof. S. Varone, directeurs de thèse

Prof. J.-F. Cordeau, rapporteur

Prof. M. G. Speranza, rapporteuse

Prof. D. Kuhn, rapporteur



ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Suisse
2017

Ce qui sauve, c'est de faire un pas. Encore un pas.
C'est toujours le même pas que l'on recommence...

— Antoine de Saint-Exupéry, 1938

Terre des Hommes

To my family and friends,
and to all who were near me during the past four years

*На моето семейство и приятели,
и на всички, които бяха до мен през изминалите четири години*

Acknowledgments

When I moved to Switzerland six years ago, my closest person here was my master's thesis supervisor Prof. Emmanuel Fragnière—a great tutor and a fantastic person whom I count among my sincerest friends. If it was not for him, I would probably never have embarked on such an arduous task. During the many coffee breaks we have had together, he would always insist that doing a PhD was the best decision in his life. Thank you Manu for giving me a push when I needed it.

Over the years that followed I had the chance to work with great professors at the HEG in Geneva. Yet, it was during my work with Prof. Sacha Varone that I first started using optimization for solving real-world practical problems. Moreover, I could experience all stages of problem solving—modeling, methodology and implementation. Thank you Sacha for giving me the opportunity to work and learn at my own pace in a relaxed and friendly atmosphere.

Starting my PhD was a reluctant decision, but at the same time strongly encouraged by Prof. Varone and Prof. Fragnière. Four years later, I regret no part of it, and I value every peak and trough I lived through. And for this, my wholehearted thanks are due to my thesis supervisor Prof. Michel Bierlaire for his immense knowledge, and for his guidance and patience. But most of all, for the personal connection through encouragement, motivation and support. I could count on him in the moments of disillusionment that every PhD student knows. Thank you Michel for the privilege to work with you, for your ability to always find a solution, and for your contagiously positive state of mind.

I want to thank my co-authors Prof. Jean-François Cordeau, Prof. Matthieu de Lapparent and Prof. Yousef Maknoon. Your knowledge and experience have been a source of inspiration. Thank you Matthieu for helping me with statistics and simulation. Thank you Yousef for the diligent proofreading of my work and for not giving up in the face of my stubborn character. Thank you both for being my friends and for the many unforgettable moments we have shared. Thank you Jean-François for hosting me in Montréal and for the chance to collaborate with you. My stay at the CIRRELT was truly transformative for my PhD experience. Your work in the field, but also your professional attitude and the sense of engagement it conveys, are something I look up to with great respect.

I wish to extend my gratitude to Prof. Corentin Fivet, Prof. Michel Bierlaire, Prof. Sacha Varone, Prof. Jean-François Cordeau, Prof. Daniel Kuhn and Prof. Maria Grazia Speranza for serving on my thesis committee. It was a real honor to have you there. Thank you for your insightful questions, for your useful comments and suggestions, but also for making my defense such a pleasant experience.

I would like to gratefully acknowledge the financial support of Switzerland's Commission for Technology and Innovation (CTI) through grant number CTI 15781.1 PFES-ES. I also wish to thank EcoWaste SA, industrial partners under the grant, for our continuous and close collaboration, for the expert advice they offered on industry and problem specific issues, and for the rich data set they provided for the experiments in this study. I would like to express my special thanks to Marc Dikötter, Phi Hung Nguyen, Gabriel Dinant and Jean-Luc Schläppi.

My sincere thanks are also due to Raphaël Lüthi and Prisca Aeby, both master's students under my supervision at the Transport and Mobility Laboratory. I take this opportunity to acknowledge their high quality work on various topics related to this thesis. I also wish to thank the editor and the anonymous referees of *Transportation Research Part B* for the constructive comments and suggestions that were instrumental in improving the quality of Chapter 2 of this thesis.

I want to thank my friends and former colleagues at the HEG for their welcoming attitude and for making my first months and years in Switzerland such a nice experience. Thank you Christophe, Tamas, Stefano, Gökcan, Bledi, Nicolas, Yannick, Francesco and many others. Thank you Moïra and Leslie for creating an atmosphere of friendliness and mutual respect in the apartments we shared. My PhD experience would not have been the same without my friends, and current and former colleagues from Transp-OR. Thank you Marianne, Anne, Mila and Emma for making sure everything runs smoothly, and for all your help for things big and small. Thank you Eva, Marija, Nikola, Anna, Stefan, Riccardo, Meri, Nicholas, Virginie, Daniel, Tomáš, Flurin, Antonin, Yuki, Aurélie, Nitish, Bilge, Matthieu, Yousef, Shadi, Jiang Hang and Amanda for all the fun and happy moments we have spent together at EPFL, in Lausanne, and beyond, and for all the fruitful scientific discussions. A big thank you to my friends and colleagues from LUTS and to all the great people I have met and continue to meet at conferences and workshops.

My experience in Lausanne has been an amazing one because of the many new friends I have made since I moved here. My warmest thanks go to Orion, Eirini, Evi, Bram, Effie, Alex, Nico, Karen, Darko, Claudia, Dan, Mihalis and Katerina. I can't close without thanking my friends and cousins from Bulgaria—Milko, Rado, Ramona, Stoyan, Misho, Elena, Svetla, Boyana, Dilyana, Vladi, Sevdalina, Metin and Pesho. I have known you for so many years, we have seen each other grow, and although we see each other less now, I feel you that much closer.

Thank you Eva for all your support! What a year it has been, with so many cherished moments, and so many other still to come. I also want to extend my boundless thanks to my family for their unconditional love—my mom Liliya, my dad Dimitar, and my sister Vladimira. They have seen me at my best and worst, and they have always been there for me. Thank you Alex for being the coolest brother-in-law. I feel like I have known you forever! Finally, I want to thank my grandparents for the best childhood I could ever have imagined.

Обичам ви!

Lausanne, November 24, 2017

Iliya Markov

Abstract

This thesis develops a unified framework for modeling and solving various classes of rich routing problems with stochastic demands, including among others the Vehicle Routing Problem (VRP) and the Inventory Routing Problem (IRP). The work is inspired by the problem of collecting recyclables from sensorized containers in the canton of Geneva, Switzerland. We start by modeling and solving the deterministic single-period version of the problem which extends the class of VRPs with intermediate facilities. It is formulated as a Mixed Integer Linear Program (MILP) which is enhanced with several valid inequalities. Due to the rich nature of the problem, general-purpose solvers are only able to tackle instances of small to medium size. To solve realistic instances, we propose a meta-heuristic approach which achieves optimality on small instances, exhibits competitive performance in comparison to state-of-the-art solution methods for special cases of the problem, and leads to important savings in the state of practice. Moreover, it highlights and quantifies the savings from allowing open tours, in which the vehicles' origin and destination depots do not coincide.

To integrate demand stochasticity, we extend the problem to an IRP over a finite planning horizon. Demand can be non-stationary and is forecast with any model that provides the expected demands and the standard deviation of the error terms, where the latter are assumed to be independent and identically distributed (iid) normal random variables. The problem is modeled as a Mixed Integer Non-Linear Program (MINLP), in which the dynamic stochastic information impacts the cost through the probability of container overflows and route failures. The solution methodology is based on Adaptive Large Neighborhood Search (ALNS) which integrates a specialized forecasting model, tested and validated on real data. The computational experiments demonstrate that our ALNS exhibits excellent performance on VRP and IRP benchmarks from the literature. The case study, which uses a set of rich IRP instances from the canton of Geneva, finds strong evidence of the added value of including stochastic information in the model. Our approach performs significantly better compared to alternative deterministic policies in its ability to limit the occurrence of overflows for the same routing cost. We also analyze the solution properties of a rolling horizon approach in terms of empirical lower and upper bounds.

This approach is generalized in a unified framework for rich routing problems with stochastic demands where we drop the assumption of iid normal error terms. We elaborate on the effects of the stochastic dimension on modeling, with a focus on the occurrence

of stock-outs/overflows and route failures, and the cost of the associated recourse actions. Tractability is achieved through the ability to pre-compute or at least partially pre-process the bulk of the stochastic information, which is possible for a general inventory policy under mild assumptions. We propose an MINLP formulation, illustrate applications to various problem classes from the literature and practice, and demonstrate that certain problems, for example facility maintenance, where breakdown probabilities accumulate over the planning horizon, can be seen through the lens of inventory routing. The case study is based on the waste collection IRP instances cited above as well as on a new set of instances for the facility maintenance problem. On the first set, we analyze the effects of our assumptions on tractability and the objective function's representation of the real cost. On the second set, we demonstrate the framework's ability to achieve the same level of occurrence of breakdowns for a significantly lower routing cost compared to alternative deterministic policies.

Key words: unified framework, rich routing problems, demand stochasticity, demand forecasting, stock-out, overflow, route failure, tractability, inventory policy, waste collection, facility maintenance, intermediate facilities, open tours, real data, Adaptive Large Neighborhood Search (ALNS), Mixer Integer Linear Program (MILP), Mixed Integer Non-Linear Program (MINLP), Vehicle Routing Problem (VRP), Inventory Routing Problem (IRP)

Résumé

Cette thèse développe un cadre unifié pour la modélisation et la résolution de diverses classes de problèmes de tournées avec demandes stochastiques, incluant entre autres le problème de tournées de véhicules (VRP) et le problème d'inventaire et de tournées (IRP). L'étude s'inspire d'un problème réel de collecte de déchets recyclables dans le canton de Genève en Suisse, dans lequel les conteneurs sont équipés de capteurs de niveau. Nous commençons par modéliser et résoudre une version déterministe du problème sur une seule période, en étendant la classe de problèmes de tournées de véhicules avec installations intermédiaires. Il est formulé en un programme linéaire mixte en nombres entiers (MILP), amélioré à l'aide d'inégalités valides. En raison de la nature riche du problème, les solveurs commerciaux ne peuvent traiter que des cas de petite à moyenne taille. Pour résoudre des cas réalistes, nous proposons une approche meta-heuristique, qui atteint l'optimum sur des instances de petite taille, présente des performances concurrentielles par rapport aux méthodes de solution de pointe sur des cas particuliers, et entraîne en pratique des économies importantes. En outre, elle met en évidence et quantifie les économies réalisées en permettant des tournées ouvertes, pour lesquelles l'origine et la destination finale de la tournée ne coïncident pas nécessairement.

Afin d'inclure l'aspect stochastique de la demande, nous étendons le problème à celui de gestion d'inventaire et de tournées (IRP) sur un horizon de planification fini. La demande considérée peut être non stationnaire et prévue à l'aide de tout modèle fournissant son espérance et l'écart-type de ses termes d'erreurs, ces derniers étant supposés indépendants et identiquement distribués (iid) selon une distribution normale. Le problème est modélisé en un programme non linéaire mixte en nombres entiers (MINLP), dans lequel l'information stochastique dynamique impacte le coût par la probabilité de débordement des conteneurs et par le risque de non validité des tournées. La méthodologie de résolution se base sur la méthode dite "Adaptive Large Neighborhood Search" (ALNS), qui intègre un modèle de prévision spécialisé, testé et validé sur des données réelles. Les tests numériques montrent que notre algorithme présente d'excellentes performances sur des instances de référence du problème de tournée de véhicules et d'inventaire dans la littérature. Le cas d'étude, qui utilise un ensemble d'instances de problèmes riches IRP du canton de Genève, valide la valeur ajoutée de l'intégration d'information stochastique dans le modèle. Notre approche est considérablement supérieure par rapport aux politiques alternatives déterministes dans sa capacité à limiter l'occurrence de débordements de conteneurs pour un coût identique de

tournées. Nous analysons également les propriétés de la solution d'une approche à horizon roulant en termes de bornes empiriques inférieures et supérieures.

Cette approche est généralisée à un cadre unifié pour les problèmes riches de tournées avec demandes stochastiques, où nous levons la restriction de termes d'erreurs supposés iid et de distribution normale. Nous analysons l'influence de la dimension stochastique sur le modèle, avec un accent mis sur le nombre d'occurrences de pénuries/débordements et de non validité de tournées, ainsi que sur le coût de leur nécessaire réaction. Le problème reste résoluble grâce à la possibilité de calculer en avance, ou au moins de partiellement traiter en amont, la majeure partie de l'information stochastique pour une politique d'inventaire générale et sous hypothèse faible. Nous proposons une formulation MINLP, illustrons par des applications de classes de problèmes variés tirés de la littérature et de la pratique, et démontrons que certains problèmes, par exemple ceux de maintenance d'installations où les probabilités de pannes s'accroissent sur l'horizon de planification, peuvent être vus dans l'optique de problèmes d'inventaires et de tournées. Le cas d'étude est basé sur les instances IRP de collecte de déchets cités ci-dessus, ainsi que sur de nouvelles instances du problème de maintenance d'installations. Sur le premier jeu de données, nous analysons les effets de nos suppositions sur la résolution et sur la représentation des coûts réels dans la fonction objectif. Sur le second jeu de données, nous démontrons la capacité du cadre unifié d'atteindre un nombre de pannes similaires pour un coût de tournées significativement moindre comparé aux politiques alternatives déterministes.

Mots-clés: cadre unifié, problèmes de tournées riches, stochasticité de la demande, prévision de la demande, pénurie, débordement, tournée non valide, résoluble, politique d'inventaire, collecte de déchets, maintenance d'installations, installations intermédiaires, tournées ouvertes, données réelles, Adaptive Large Neighborhood Search (ALNS), programme linéaire mixte en nombres entiers (MILP), programme non linéaire mixte en nombres entiers (MINLP), problème de tournées de véhicules (VRP), problème d'inventaire et de tournées (IRP).

Contents

Acknowledgments	i
Abstract (English/Français)	v
List of Tables	xiii
List of Figures	xv
List of Algorithms	xvii
1 Introduction	1
1.1 Context, State of the Art and Motivation	1
1.2 Objectives	7
1.3 Contributions	8
1.4 Thesis Organization	10
2 The Waste Collection VRP	11
2.1 Related Literature	12
2.2 Exact Approach	14
2.2.1 Formulation	14
2.2.2 Variable Fixing and Valid Inequalities	18
2.3 Meta-heuristic Approach	20
2.3.1 Feasibility	20
2.3.2 Initial Solution Construction	22
2.3.3 Multiple Neighborhood Search	23
2.4 Numerical Experiments	27
2.4.1 Evaluation on Small Instances	27
2.4.2 Tests on Benchmark Instances from the Literature	31
2.4.3 Case Study	34
2.5 Summary	36
3 The Waste Collection IRP with Stochastic Demands	37
3.1 Related Literature	38
3.2 Formulation	41
3.2.1 Forecasting Model	42

Contents

3.2.2	Stochastic IRP Model	43
3.2.3	Model Reformulations for Solving Benchmark Instances	50
3.3	Adaptive Large Neighborhood Search	52
3.3.1	Solution Representation	53
3.3.2	Operators	55
3.3.3	Algorithmic Modifications for Solving Benchmark Instances	58
3.4	Numerical Experiments	59
3.4.1	Parameter Tuning	60
3.4.2	Benchmark Results	61
3.4.3	Case Study	64
3.5	Summary	75
4	A Unified Framework for Rich Routing Problems with Stochastic Demands	77
4.1	Related Literature	78
4.1.1	Rich Vehicle and Inventory Routing Problems	78
4.1.2	Health Care Routing Problems	79
4.1.3	Waste Collection Routing Problems	79
4.1.4	Maritime Routing Problems	80
4.1.5	Discussion	81
4.2	Key Concepts and Modeling Elements	81
4.3	Capturing Demand Stochasticity	85
4.3.1	Demand Decomposition and Forecasting	85
4.3.2	Demand Point Probabilities	85
4.3.3	Route Failure Probabilities	89
4.4	Optimization Model	91
4.4.1	Objective Function	92
4.4.2	Deterministic Constraints	95
4.4.3	Probabilistic Constraints	97
4.5	Application Examples	97
4.5.1	The Vehicle Routing Problem	98
4.5.2	The Health Care Inventory Routing Problem	98
4.5.3	The Waste Collection Inventory Routing Problem	98
4.5.4	The Maritime Inventory Routing Problem	99
4.5.5	The Facility Maintenance Problem	99
4.6	Numerical Experiments	101
4.6.1	Instances	101
4.6.2	Solving the Waste Collection Inventory Routing Problem	102
4.6.3	Solving the Facility Maintenance Problem	105
4.7	Summary	108

5 Conclusion	111
5.1 Main Findings	111
5.2 Practical Implications	113
5.3 Future Research Directions	114
A Waste Collection IRP: Additional Analysis	117
A.1 Waste Collection IRP: Collection Strategies	117
A.2 Waste Collection IRP: Effect of Lower Truck Capacity on the Solution Cost . . .	119
A.3 Waste Collection IRP: Effect of Open Tours on the Solution Cost	120
B Equivalence of Stock-out and Overflow Probabilities	123
C Processing Stochastic Information	125
C.1 Pre-computing the Stock-out Probabilities	125
C.2 Partially Pre-processing the Route Failure Probabilities	126
Bibliography	127
Curriculum Vitae	141

List of Tables

1.1	Structural Classification for the IRP (Coelho et al., 2014b)	5
2.1	Notations	16
2.2	Algorithmic Parameters	27
2.3	MNS vs Solver on Modified Schneider et al. (2014) Instances	29
2.4	Evaluation of MNS Neighborhoods and Their Combinations	30
2.5	Comparison Against the BKS to the MDVRPI (Crevier et al., 2007) Instances	31
2.6	Savings from Allowing Open Tours in the MDVRPI (Crevier et al., 2007) Instances	32
2.7	Savings from Home Depot Optimization and Allowing Open Tours in the MDVRPI (Crevier et al., 2007) Instances	33
3.1	Notations	41
3.2	Algorithmic Parameters	60
3.3	Results on Archetti et al. (2007) Instances	62
3.4	Results on Crevier et al. (2007) Instances	63
3.5	Results on Taillard (1999) Instances	64
3.6	Probabilistic Policies: Basic Results for Cost Analysis on Real Data Instances	66
3.7	Probabilistic Policies: Key Performance Indicators for Cost Analysis on Real Data Instances	66
3.8	Probabilistic Policies: Container Overflows and Route Failures for Real Data Instances	67
3.9	Driving Factors for the Occurrence of Container Overflows	69
3.10	Alternative Policies: Basic Results for Cost Analysis on Real Data Instances	70
3.11	Alternative Policies: Key Performance Indicators for Cost Analysis on Real Data Instances	71
3.12	Alternative Policies: Container Overflows and Route Failures for Real Data Instances	71
3.13	Analysis of Rolling Horizon DSIRP Bounds	75
4.1	Notations	84
4.2	Impact of Empirical Distribution Functions on Tractability	103
4.3	Basic Results for Model (FMP1) vs. Model (FMP2)	106
4.4	Performance Indicators for Model (FMP1) vs. Model (FMP2)	107

List of Tables

4.5	Performance Indicators for Model (FMPD)	108
A.1	Average Level of All Containers by Day	117
A.2	Average Level of Collected Containers by Day	118
A.3	Number of Instances with Container Collections by Day	118

List of Figures

1.1	Tour Example	3
2.1	Multiple Neighborhood Search Operators	24
2.2	Geneva Service Area	34
2.3	Distance Improvements Compared to the Currently Used Software	35
3.1	Container State Probability Tree	44
3.2	Feasibility Graph of the Reorder Dumps Operator	58
3.3	Average Cost of Overflows at Different Percentiles of the Simulated Scenarios	68
3.4	Overflows for All Instances at Different Percentiles of the Simulated Scenarios	68
3.5	Comparison of Routing Cost for Probabilistic and Alternative Policies	72
3.6	Comparison of Container Overflows and Route Failures for Probabilistic and Alternative Policies	73
4.1	Discrete Maximum Level Policy Example	83
4.2	Demand Point State Probability Tree	88
4.3	Breakdown Probabilities for Different Values of α	102
4.4	Objective Function's Overestimation of the Real Cost	105
4.5	Comparison of Routing Cost and Breakdowns for Model (FMP1) vs. Model (FMP2)	108
A.1	Effect of Lower Truck Capacity on the Solution Cost	119
A.2	Effect of Open Tours on the Solution Cost	120

List of Algorithms

2.1	Temporal Feasibility Algorithm	21
2.2	Multiple Neighborhood Search	25
3.1	Adaptive Large Neighborhood Search	53
4.1	Construction of the Set of Supply Point Delimited Trips \mathfrak{S}_k for Vehicle k	90

1 Introduction

This chapter borrows from the articles:

Markov, I., Varone, S., and Bierlaire, M. (2016). Integrating a heterogeneous fixed fleet and a flexible assignment of destination depots in the waste collection VRP with intermediate facilities, *Transportation Research Part B: Methodological* 84:256-273.

Markov, I., Bierlaire, M., Cordeau, J.-F., Maknoon, Y., and Varone, S. (2016). Inventory routing with non-stationary stochastic demands. Technical report TRANSP-OR 160825. Transport and Mobility Laboratory, EPFL, Lausanne, Switzerland.

Markov, I., Bierlaire, M., Cordeau, J.-F., Maknoon, Y., and Varone, S. (2017). A general framework for routing problems with stochastic demands. *Proceedings of the 17th Swiss Transport Research Conference (STRC)*, May, 17-19, 2017.

The work therein has been performed by the author in collaboration with Prof. Michel Bierlaire, Prof. Jean-François Cordeau, Prof. Yousef Maknoon and Prof. Sacha Varone.

Section 1.1 of this chapter introduces the context and provides a focused analysis of the literature pertinent to this research, motivating the work undertaken here. This is followed by Section 1.2 which identifies the objectives we set to achieve and Section 1.3 which lists the contributions we make to the state of the art. Finally, Section 1.4 outlines the organization of the thesis chapter by chapter.

1.1 Context, State of the Art and Motivation


The Vehicle Routing Problem (VRP) is an integer programming and combinatorial optimization problem that seeks to find the cheapest set of tours that a fleet of vehicles should perform to serve a set of customers. In its basic form, the VRP considers a single depot where

each tour starts and ends, and a homogeneous fleet of capacitated vehicles. Each customer has a fixed demand of a single commodity and the number of customers in each tour is only limited by the vehicle capacity. The VRP was formally introduced in the seminal work of Dantzig and Ramser (1959) in the context of fuel delivery and is one of the most practically relevant and widely studied problems in operations research. A generalization of the VRP, the Inventory Routing Problem (IRP) introduces a planning horizon and seeks to optimize simultaneously the vehicle tours, delivery times and delivery quantities. The seminal work on the IRP was motivated by the delivery and inventory management of industrial gases (Bell et al., 1983). The literature on the VRP, the IRP and their many variants is vast, driven both by their mathematical properties from an optimization point of view, but also due to their numerous practical applications in the distributions and collection of goods and the transportation of people. The need to solve larger and richer routing problems has pushed researchers over the past decades to develop sophisticated and efficient solution methodologies.

Rich routing problems are multi-constrained routing problems with a variety of real-world features. In such a setting, the fleet can be heterogeneous instead of homogeneous, and each vehicle can perform multiple tours per day, instead of one, and visit multiple customers and replenishment facilities, subject to time windows and accessibility restrictions. Depending on the application, there could be multiple depots with the possibility of open tours that have different origin and destination depots, or multi-day tours that last over several days. Drivers are subject to regulations on maximum working hours, while equity considerations might imply that all drivers work similar hours. Customers may have preferences for a given driver and visit periodicity. Because of their inherent difficulty, such problems have seen increased academic interest in recent years due to methodological and technological progress (Lahyani et al., 2015). Another defining characteristic of real-world problems is uncertainty, which can present itself in a variety of ways, such as uncertain demand quantities, uncertain customer presence, uncertain travel and service times, etc (Gendreau et al., 2016). The rich routing features, combined with the necessity of tracking inventory over the planning horizon in the case of the IRP, inevitably compound the effects of uncertainty. Failure to account for uncertainty often leads to solutions that are highly suboptimal or even infeasible given the realizations of the uncertain parameters (Louveaux, 1998).

This thesis develops a unified framework for modeling and solving various classes of rich routing problems with stochastic demands, including among others the VRP and the IRP. The work is inspired by the problem of collecting recyclables from sensorized containers in the canton of Geneva, Switzerland. In this context, we solve a rich IRP in which a heterogeneous vehicle fleet collects recyclable waste over a planning horizon of approximately one week. Waste containers and collection vehicles are flow-specific, hence the problem can be solved separately for each waste flow. As shown in Figure 1.1, each tour starts and ends at a depot, not necessarily the same, and is a sequence of collections followed by disposals at the available dumps, with a mandatory visit to a dump before the end of the tour. Dumps are recycling plants, and there could be multiple dumps for the collected waste flow which can

Figure 1.1: Tour Example



be used when and as needed along the tour. We consider time windows for the depots, containers and dumps. A tour is also limited by the legal duration of the working day. Site dependencies, or accessibility restrictions, apply to certain points, for example to containers located in narrow streets that cannot be accessed by big collector trucks. With container capacities in the range of one to three thousand liters, the infrastructure is relatively sparse with collection points typically several hundred meters apart. As a consequence, we are faced with a node-routing problem, as opposed to the arc-routing problems applicable to household waste collection (Golden et al., 2002) where demands along the street segments can be clustered.

Containers are equipped with ultrasound sensors and transmit their levels to a central database at the start of each day. Using the historical data, a forecasting model estimates the expected daily demands and a measure of uncertainty associated with them. This information is used for calculating the probabilities of container overflows and route failures during the planning horizon, the two types of undesirable events that we consider. The concept of overflow is self-explanatory, while a route failure refers to an event where the vehicle runs out of capacity before the next scheduled dump visit due to higher than expected demand realizations (Dror and Trudeau, 1986). The IRP literature typically uses simple forecasting techniques, if at all. Our framework allows the use of any state-of-the-art approach that provides the expected demands and a measure of uncertainty associated with them in the form of a probability distribution.

Over a single day, the described problem is an extension of the Vehicle Routing Problem with Intermediate Facilities (VRP-IF), in which the facilities represent the dumps where vehicles stop to empty. Although the research on waste collection VRP variants spans several decades, the VRP-IF itself has been relatively less studied. One of its first applications to waste collection is that of Beltrami and Bodin (1974) who apply it to the case of New York City. More recently, applications to waste collection have appeared in Angelelli and Speranza (2002a,b), Kim et al. (2006), Benjamin (2011), Hemmelmayr et al. (2013) and others. The VRP-IF has wide practical applicability outside waste collection too, as evidenced by the variety of solution techniques developed for it, which include meta-heuristic (Tarantilis et al., 2008), hybrid (Crevier et al., 2007) and fully exact methods (Muter et al., 2014).

Yet, despite the practical relevance of these types of problems, many important realistic features are often omitted. For example, the literature on the VRP-IF often assumes homogeneous fleets, whereas in industry fleets either start as heterogeneous or become such as new vehicles are added or old ones are replaced. Many studies have treated heterogeneous fleets in an ad-hoc manner but Taillard (1999) is the first to formulate the heterogeneous fixed fleet VRP, a generalization of the vehicle fleet mix problem with a fixed number of vehicles of each type. The years that followed saw a strong scientific interest in this problem and the development of state-of-the-art heuristic (Subramanian et al., 2012; Penna et al., 2013) and exact methodologies (Baldacci and Mingozzi, 2009). Our work considers intermediate facilities and a heterogeneous fixed fleet simultaneously, resulting in a more complex problem where the cost attractiveness of smaller vehicles is counterbalanced by the need for more frequent dump visits.

As illustrated in Figure 1.1, our setup also considers open tours (Sariklis and Powell, 2000; Repoussis et al., 2007, 2010; Yousefikhoshbakht et al., 2014), with origin and destination depots not necessarily the same. The possibility of open tours relies on the presence of multiple depots as well (Cordeau et al., 1997, 2001; Dondo and Cerdá, 2006, 2007, 2009; Baldacci and Mingozzi, 2009; Bettinelli et al., 2011). Furthermore, the addition of constraints on time windows and maximum tour duration (Solomon, 1987; Savelsbergh, 1992; Cordeau et al., 2001; Cordeau and Laporte, 2001; Cordeau et al., 2004; Polacek et al., 2004), and site dependencies (Nag et al., 1988; Chao et al., 1999; Cordeau and Laporte, 2001) results in a rich routing problem. Being generalizations of the basic VRP, such problems are *NP*-hard and are known to be notoriously difficult to solve. Not surprisingly, they have been the subject of increased academic interest in the last decade due to the methodological and technological progress (Lahyani et al., 2015), enabling the solution of larger and richer problems. A case in point is the work of Pisinger and Ropke (2007), who use the Adaptive Large Neighborhood Search (ALNS) algorithm of Ropke and Pisinger (2006a) to solve several classes of vehicle routing problems, including the VRP with time windows, the multi-depot VRP, the site-dependent VRP, and the open VRP. Pisinger and Ropke (2007) test the ALNS on 486 benchmark instances from the literature and improve 183 best known solutions. In recent years, advances have also led to the appearance of a number of unified frameworks for different classes of routing problems (Ropke and Pisinger, 2006b; Pisinger and Ropke, 2007; Baldacci and Mingozzi, 2009; Vidal et al., 2014; Koç, 2016; Kritzing et al., 2017). The success of the ALNS on many classes of routing problems over the past decade motivated the development of our own implementation, specifically designed to exploit the structure of our framework.

A further source of complexity arises from the presence of stochastic elements (Gendreau et al., 2016), in our particular case stochastic demands. While sensor information provides the current container levels, future ones are only available as forecasts. Over a multi-day planning horizon, the problem becomes an IRP, which determines simultaneously the vehicle tours and visit days, and as a consequence the collection quantities. Coelho et al. (2014b) carry out a detailed survey of the IRP literature during the past thirty years. Table 1.1 describes

Table 1.1: Structural Classification for the IRP (Coelho et al., 2014b)

Criterion	Classification
Time horizon	Finite (rolling)
Structure	Many-to-many
Routing	Multiple
Inventory policy	Order-Up-to (OU)
Inventory decisions	Back-ordering (with a penalty and limit)
Fleet composition	Heterogeneous
Fleet size	Multiple (fixed)

our problem in terms of the structural classification scheme they propose. Specifically, it is defined over a finite planning horizon and solved in a rolling horizon fashion. That is, when the IRP is solved, only the decisions on the first day are implemented, after which the horizon is shifted by a day and the problem is re-solved with updated container levels and forecasts. This approach has been central to the IRP since the seminal works in this field (e.g. Bard et al., 1998b). To the contrary, solving the problem day by day in isolation would lead to myopic decisions, often or always postponing collections for the future so as to minimize the routing cost today (Trudeau and Dror, 1992). Our problem also considers multiple containers that are emptied into multiple dumps (intermediate facilities), which identifies the structure as many-to-many. While intermediate facilities have been considered in the IRP literature, the simultaneous presence of the variety of rich routing features present in our problem has rarely been systematically treated. Routing involves multiple containers at a time and when a container is visited it is fully emptied, i.e. according to an Order-Up-to (OU) level inventory policy (Bertazzi et al., 2002). Experience suggests that overflowing containers continue serving demand because people place the waste beside them. As a consequence, container overflows are served at a penalty (back-order) but the number of back-order days is limited to one. The fleet is fixed and heterogeneous, with vehicles possibly having different cost and capacity characteristics. Information-wise, we have stochastic demands and knowledge about their distributions. The rolling horizon approach introduces dynamism where new container level and forecasting information is revealed each day.

There is a variety of modeling approaches for stochastic optimization problems, of which routing problems with stochastic demands are a special case. Scenario generation and stochastic modeling based on Markov decision processes both lead to problems that suffer from the curse of dimensionality for realistic-size instances (Pillac et al., 2013). Approximate dynamic programming (Powell, 2011) helps alleviate the problem in the latter case. In their recent work, Rossi et al. (2017) also note the instance size limitations of dynamic programming in solving the bowser routing problem, a special version of the IRP, and propose heuristic approximations. Robust optimization maintains feasibility for a given budget of uncertainty and is distribution-free. It relies on specific reformulations depending on whether parameter uncertainty in the standard-form mathematical program appears column-wise (Soyster,

1973), row-wise (Bertsimas and Sim, 2003, 2004), or only in the right-hand side (Minoux, 2009). Yet, complications arise if there is inter-row dependency in the uncertainty on the right-hand side (see Delage and Iancu, 2015). We do not see this approach very often used for routing problems (Gendreau et al., 2016), but we should mention the works of Sungur et al. (2008) and Gounaris et al. (2013) who treat stochastic demands for the VRP and Aghezzaf (2008) and Solyalı et al. (2012) for the IRP. Chance constrained approaches guarantee that a constraint will be satisfied with a given probability. These are appropriate if uncertainty appears row-wise and have typically been used to model route failures in vehicle routing problems with stochastic demands (see references in Gendreau et al., 2014).

The use of a particular approach has a strong influence on how the problem at hand is being viewed. Both robust optimization and chance constraints are risk-oriented approaches, shifting the treatment of uncertainty to the constraints. They also leave open the question of how to set the budget of uncertainty or the distribution percentile for the chance. Robust optimization in particular is less relevant for our problem where container overflows and route failures are not disastrous events. Their states are frequently revisited, unlike what is usually the case in robust optimization. Furthermore, container overflows and route failures have a monetary expression which should figure in the total expected cost incurred by the collector. Thus, the integration of probability information in the objective is used to provide a monetary dimension to these stochastic events, and their associated recourse actions, resulting in a cost-oriented approach. Indeed, it often pays off taking a small risk if other cost components can be significantly reduced as a consequence. Scenario generation approaches capture the cost of recourse but would be very cumbersome computationally for a rich routing problem like ours. Chance constraints, on the other hand, can be easily integrated in our framework. Lastly, we highlight that the vast majority of the distribution-based approaches in the literature on routing problems assume independent and identically distributed (iid) demands from the normal distribution (see Gendreau et al. (2016) for the case of the stochastic VRP).

The waste collection problem described here is a rich IRP with stochastic demands, and while our modeling techniques are conceived within this context, they are extensible to a number of other problems as well, ranging from deterministic VRPs over a single period all the way to maritime IRPs with stochastic demands. For all practical purposes, distribution and collection problems are identical, since distribution can be viewed as the collection of empty space, and vice versa. Overflows also have their counterpart in a distribution context in the form of stock-outs. The concept of route failure is universal as well. Moreover, there need not be any restriction to routing problems that consider demand, as the probability of undesirable events not related to demand can be treated in the same or similar ways. A case in point is the facility maintenance problem, in which the probability of a facility breakdown is a function of the number of days that have passed since its most recent visit. In other words, the probability of breakdown accumulates in a way similar to inventory, thus allowing us to use the same ideas and modeling techniques to solve this problem. Bringing these elements and observations together, this thesis culminates in a unified framework for

modeling and solving rich routing problems in the presence of non-stationary stochastic demands, and a powerful solution methodology, based on ALNS, that is able to handle the resulting complexity.

Starting from the waste collection problem, we gradually build our unified framework and illustrate how it adapts to the original problem, but also to a number of vehicle, inventory and other routing problems from the literature and practice. We discuss the generality of our approach in terms of the richness of its routing features and the probability distributions that it can handle with fully or partially relaxed iid normal assumption. The latter addresses the gap between theory and practice when it comes to stochastic routing problems (Gendreau et al., 2016), and is confirmed with a successful application to rich problem instances derived from real data. Still, we carefully list all the assumptions and restrictions that we impose in order to keep the approach tractable. Tractability enters the modeling framework at the design level in the ability to pre-compute or at least partially pre-process the bulk of the stochastic information, which is not only possible for the OU inventory policy mentioned above, but also for a generalized one under mild assumptions. In sum, while the practical motivation behind this work is the solution of a real-world waste collection problem with important economic consequences for any city, there is also a more abstract theoretical motivation. We are well aware of the pitfalls of a one-size-fits-all approach and are far from proposing one. Nevertheless, we are convinced that the ideas we develop and present here, and which are inspired from the work on the waste collection problem, are general enough and extensible even further.

1.2 Objectives

In the process of developing the unified framework described in Section 1.1, this thesis sets the following objectives:

1. *Integration of rich routing features.* The modeling and solution approach needs to be applicable to real-world problems, which are characterized by complex physical and temporal constraints.
2. *Integration of demand forecasting.* The future is uncertain but forecasting techniques or expert knowledge are crucial in isolating trends. Forecasts are not only useful in making decisions for the future but even more so in shaping the decisions we make for today.
3. *Integration of demand uncertainty through explicit modeling of undesirable events and their associated recourse actions.* Stochastic demand is a source of risks whose nature depends on the problem at hand. We consider two principal risks common to routing problems: 1) undesirable events at the demand points, e.g. overflows or other, depending on application, and 2) route failures. We model explicitly the probabilities of these events, their costs and the costs of their associated recourse actions. Thus, we

model the resulting cost of demand uncertainty that one has to pay.

4. *Tractability.* Routing problems are typically operational short-term problems. As such, they need to be solved quickly and efficiently. To do this, we need to consider tractability at the design phase through the manner in which stochasticity is captured and treated and at the solution phase in the form of efficient algorithms. Our goal is to be able to solve realistic-size instances in the order of minutes or tens of minutes. The difficulty of our problem limits the use of fully exact approaches in favor of meta-heuristics.
5. *Generality and genericity.* The problem we start from is a rich waste collection routing problem, but the concepts, models and algorithms we develop, and ideally the conclusions we draw, should be extensible to vehicle, inventory and other routing problems from different application areas. Likewise, the solution methodology should be powerful but remain generic enough and modular to allow the integration of new routing features. The ALNS approach is well-suited for this purpose.
6. *Successful application to a real case study.* Real-world applications pose a great challenge but at the same time provide the most meaningful evaluation of models and algorithms for rich routing problems. We work with a rich data set coming from the canton of Geneva, Switzerland, which includes the city of Geneva and the surrounding area, home to approximately half a million people and a dynamic financial and diplomatic center with millions of visitors annually.

1.3 Contributions

In achieving the objectives outlined in Section 1.2, this thesis makes the fundamental and practical contributions described below.

The vehicle routing problem with intermediate facilities (Chapter 2):

- The integration of a heterogeneous fixed fleet into the VRP-IF, with multiple depots and the possibility of open tours with different origin and destination depots.
- The development of a Mixed Integer Linear Program (MILP), which is enhanced with several valid inequalities with significant impact on computation time.
- The development of a meta-heuristic approach with multiple neighborhood search, which achieves optimality on small instances and exhibits competitive performance in comparison to state-of-the-art solution methods for special cases of our problem. It also confirms the benefit of open tours and the potential for important financial savings in the current state of practice in the canton of Geneva.

The waste collection inventory routing problem with stochastic demands (Chapter 3):

- The utilization of a purpose-designed demand forecasting model, tested and validated on real data, which assumes iid demand error terms.
- The explicit modeling of undesirable events and their associated recourse actions by incorporating dynamic probabilistic information in the objective function.
- The integration of a variety of rich routing features traditionally absent or rarely considered in the IRP literature, such as a heterogeneous fixed fleet, intermediate facilities, time windows, maximum tour duration, accessibility restrictions, etc.
- The extensive computational testing of our ALNS algorithm, which demonstrates its excellent performance on IRP and VRP benchmarks from the literature. We evaluate the benefit of integrating uncertainty in the decision process, with our approach significantly outperforming alternative deterministic policies in limiting the occurrence of container overflows for the same routing cost. We also analyze the solution properties of a rolling horizon approach for a dynamic and stochastic version of the problem and derive empirical lower and upper bounds on its solution cost.

The unified framework for rich routing problems with stochastic demands (Chapter 4):

- The generalization of the forecasting model through the complete or partial relaxation of the assumption on iid normal error terms.
- The generalization of the undesirable events and the proof of equivalence between distribution and collection problems.
- The preservation of computational tractability in the face of the above generalizations through the ability to pre-compute or at least partially pre-process the bulk of the stochastic information for a general inventory policy. This comes at the expense of several assumptions and simplifications whose effect on the solution cost is shown to be marginal through a simulation-validation approach.
- The generality and practical relevance of the approach. We integrate the probabilistic information into a Mixed Integer Non-Linear Program (MINLP), illustrate applications to various problem classes from the literature and practice, such as health care, waste collection, and maritime inventory routing, and demonstrate that problems like facility maintenance can be seen through the lens of inventory routing. Extending the ALNS, the computational experiments focus on the topic of complexity vs. tractability. In addition, for realistic instances of the facility maintenance problem, we show the framework's ability to achieve the same level of occurrence of breakdowns for a significantly lower routing cost compared to alternative deterministic policies.

1.4 Thesis Organization

The rest of this thesis is organized as follows. Chapter 2 models the deterministic single-day waste collection problem—an extension of the VRP with intermediate facilities. The chapter is based on the following article:

- Markov, I., Varone, S., and Bierlaire, M. (2016). Integrating a heterogeneous fixed fleet and a flexible assignment of destination depots in the waste collection VRP with intermediate facilities. *Transportation Research Part B: Methodological*, 84: 256–273.

Preliminary ideas and results are also discussed in:

- Markov, I., Varone, S., and Bierlaire, M. (2014). Vehicle routing for a complex waste collection problem. *Proceedings of the 14th Swiss Transport Research Conference (STRC)*, May 14-16, 2014, Ascona, Switzerland.
- Markov, I., Varone, S., and Bierlaire, M. (2015). The waste collection VRP with intermediate facilities, a heterogeneous fixed fleet and a flexible assignment of origin and destination depot. Technical report TRANSP-OR 150212, Transport and Mobility Laboratory, EPFL, Lausanne, Switzerland.

Chapter 3 integrates demand stochasticity and extends the problem to an IRP over a finite planning horizon. The chapter is an extension of the following technical report:

- Markov, I., Bierlaire, M., Cordeau, J.-F., Maknoon, Y., and Varone, S. (2016). Inventory routing with non-stationary stochastic demands. Technical report TRANSP-OR 160825, Transport and Mobility Laboratory, EPFL, Lausanne, Switzerland.

Preliminary ideas and results, in particular those related to demand forecasting, are also discussed in:

- Markov, I., Lapparent, M. (de), Bierlaire, M., and Varone, S. (2015). Modeling a waste disposal process via a discrete mixture of count data models. *Proceedings of the 15th Swiss Transport Research Conference (STRC)*, April 15-17, 2015, Ascona, Switzerland.

Chapter 4 generalizes the approach in a unified framework for rich routing problems with stochastic demands. The chapter extends the ideas presented in:

- Markov, I., Bierlaire, M., Cordeau, J.-F., Maknoon, Y., and Varone, S. (2017). A general framework for routing problems with stochastic demands. *Proceedings of the 17th Swiss Transport Research Conference (STRC)*, May, 17-19, 2017.

Finally, Chapter 5 closes with a summary of the main findings and conclusions. It discusses the practical implications of this work and identifies promising and pertinent areas of future research.

2 The Waste Collection VRP

This chapter is based on the article:

Markov, I., Varone, S., and Bierlaire, M. (2016). Integrating a heterogeneous fixed fleet and a flexible assignment of destination depots in the waste collection VRP with intermediate facilities, *Transportation Research Part B: Methodological* 84:256-273.

The work therein has been performed by the author in collaboration with Prof. Michel Bierlaire and Prof. Sacha Varone.

This chapter models and solves the single-day waste collection VRP, an extension of the class of VRPs with intermediate facilities. Given a daily time discretization and sensor information at the beginning of the day, the problem is deterministic. This rich VRP variant includes a heterogeneous fixed fleet of capacitated vehicles collecting recyclable waste from a set of containers. There is also a set of recycling facilities, or dumps, where vehicles stop to empty the collected waste when and as needed along the tour. There is no limit on the number of dump visits and there is a mandatory dump visit before the end of the tour, i.e. a tour terminates with an empty vehicle.

We have multiple depots and the possibility of open tours with different origin and destination depots. The realistic setting also includes time windows and a maximum tour duration. Additionally, there is a mandatory break after a predefined number of hours of continuous work. Accessibility restrictions limit the types of vehicles that can visit certain containers. For example, big collector trucks may not be able to reach containers in narrow streets. The objective function captures the principal cost components faced by a typical firm, i.e. deployment cost, travel distance and travel time related cost. It also includes a special cost term capturing the potential relocation cost in the future due to open tours. Given the difficulty of this problem, we develop an exact approach based on a mathematical formulation enhanced with valid inequalities which can solve small to medium-size instances. To solve realistic instances, we propose a meta-heuristic approach based on multiple neighborhood search.

The chapter is organized as follows. Section 2.1 is a brief analysis of the related literature. Sections 2.2 and 2.3 present the exact approach and the meta-heuristic approach, respectively. Section 2.4 discusses the results of the numerical experiments, and Section 2.5 ends with a summary of the main findings and contributions.

2.1 Related Literature

Bard et al. (1998a) propose the first exact approach, based on branch-and-cut, for the VRP with Intermediate Facilities (VRP-IF) in a distribution context. With it, they are able to solve to optimality certain instances with up to 20 customer nodes. Bard et al. (1998b) extend this setup to an IRP framework which they solve in a rolling horizon fashion. Angelelli and Speranza (2002b) apply a modification of Cordeau et al.'s (1997) unified Tabu Search (TS) algorithm to a Periodic VRP-IF (PVRP-IF) with features such as service durations and a maximum tour duration. In Angelelli and Speranza (2002a), this framework is used to analyze the operational cost benefits of different waste collection policies in Val Trompia, Italy and Antwerp, Belgium.

Kim et al. (2006) include time windows and a driver break in the waste collection VRP-IF, explicitly considering also features such as tour compactness and workload balancing. Their solution approach, an extension of Solomon's (1987) insertion algorithm followed by simulated annealing, leads to a significant reduction in the number of tours and substantial financial savings at a major US waste collection company (see Sahoo et al., 2005). Kim et al. (2006) are also the first to propose a set of 10 benchmark instances for the VRP-IF, involving up to 2092 stops and 19 intermediate disposal facilities. The multi-objective genetic algorithm of Ombuki-Berman et al. (2007), the variable neighborhood tabu search of Benjamin (2011) and the ALNS of Buhrkal et al. (2012) are also applied on these instances, leading to distance improvements of 10-15% and using fewer vehicles. Buhrkal et al.'s (2012) approach also leads to a distance improvement of 30-45% at a Danish waste collection company.

Crevier et al. (2007) propose the Multi-Depot VRP with Inter-depot routes (MDVRPI). Although the setup is closely related, it was originally applied in a distribution context. The MDVRPI is non-periodic, no time windows or driver breaks are considered and, in the general case, depots and intermediate facilities coincide. Crevier et al. (2007) use the Adaptive Memory (AM) principle of Rochat and Taillard (1995) and decompose the problem into multi-depot, single-depot and inter-depot subproblems which are solved using Cordeau et al.'s (1997) TS. A solution to the MDVRPI is obtained through a set covering formulation and improved by a modified version of the TS.

Crevier et al. (2007) create two sets of MDVRPI instances with 48 to 288 customers and a fixed homogeneous fleet stationed at one depot, with the rest of the depots acting only as intermediate facilities. These instances are used by Tarantilis et al. (2008) and Hemmelmayr et al. (2013) who propose, respectively, a hybrid guided local search and a Variable Neighborhood

Search (VNS) with a dynamic programming procedure for the insertion of the intermediate facilities in the tours. Both articles report improvements over the results of Crevier et al. (2007) with computation times close to one hour for the largest instances. Muter et al. (2014) develop a branch-and-price algorithm for the MDVRPI and solve several sets of instances derived from those of Crevier et al. (2007) to optimality. Only Hemmelmayr et al. (2013) apply their methodology to a PVRP-IF faced by a waste collection company and achieve a 25% reduction in the routing cost. Hemmelmayr et al. (2014) combine this problem with the bin allocation problem and study the cost trade-off between less frequent visits and larger bin sizes. They develop a matheuristic approach with a VNS for the routing problem and a mathematical model for the bin allocation problem, and compare a hierarchical and an integrated approach.

Another related problem class is the routing of electric or alternative fuel vehicles, where we have recharging or refueling decisions in lieu of emptying decisions. Conrad and Figliozzi (2011) consider the Recharging VRP (RVRP), where electric vehicles can recharge at customer locations with time windows. Erdoğan and Miller-Hooks (2012) consider the Green VRP (G-VRP), where vehicles use a sparse alternative fuel infrastructure. Results on medium-size random instances show that spatial characteristics have a significant impact on the optimality gap, which appears to be related to the number of facilities. Larger instances are used to analyze the effects of increasing the number of customers, facility availability and driving range limits.

Schneider et al. (2014) solve the Electric VRP with Time Windows and recharging stations (E-VRPTW). The problem features variable recharging times based on remaining battery charge and a hierarchical objective function minimizing number of vehicles first and travel distance second. The proposed hybrid VNS/TS improves the results of Erdoğan and Miller-Hooks (2012) by 8-15% and obtains competitive results on the MDVRPI sets of Crevier et al. (2007) and Tarantilis et al. (2008). Schneider et al. (2015) combine recharging and reloading facilities in the VRP with Intermediate Stops (VRPIS). Contrary to the E-VRPTW, here the objective function is weighted rather than hierarchical. The authors propose an ALNS, which is able to match or improve the results of Schneider et al. (2014) on the G-VRP instances at a fraction of the computation time. Convincing results are also obtained for the MDVRPI instances of Crevier et al. (2007) and Tarantilis et al. (2008). Adler and Mirchandani (2017) propose a branch-and-price algorithm and a heuristic for a multi-depot scheduling problem for alternative fuel vehicles. Recent surveys of the relevant literature are available in Moghaddam (2015) and Pelletier et al. (2016).

Regarding the vehicle fleets, Kim et al. (2006) and the related papers on the VRP-IF assume an unlimited homogeneous fleet. The PVRP-IF of Angelelli and Speranza (2002b), the MDVRPI, RVRP, G-VRP, E-VRPTW and VRPIS also assume a homogeneous fleet, albeit limited. More recently, Sassi et al. (2014), Goeke and Schneider (2015), Mancini (2016), and Hiermann et al. (2016) have started filling the gap by considering conventional and alternative fuel vehicles simultaneously. Yavuz and Çapar (2017) study the adoption of alternative fuel vehicles

into gasoline and diesel fuel fleets, considering various objective functions, and discuss performance and managerial implications. Taillard (1999) was the first to formally define the Heterogeneous Fixed Fleet VRP (HFFVRP). Being a generalization of the Vehicle Fleet Mix Problem (VFMP), the HFFVRP is *NP*-hard and more difficult than the classical VRP or the VFMP. Taillard's (1999) solution approach relies on heuristic column generation with AM, and vehicle assignment costs are calculated at each iteration. He adapts the eight largest VFMP instances of Golden et al. (1984) to the HFFVRP by specifying the number of vehicles of each type and their variable costs. The best heuristic approaches on these benchmarks are due to Penna et al. (2013) and Subramanian et al. (2012), the latter also being the fastest. The only fully exact method is that of Baldacci and Mingozzi (2009). They prove the optimality of seven of the best known solutions to the instances with variable costs only, and six in the case where both fixed and variable costs are considered.

The originality of our problem is thus reinforced by the general lack of literature treating the heterogeneous fixed fleet VRP-IF despite its wide practical application. The combined presence of a heterogeneous fixed fleet and intermediate facilities results in a more complex problem where the cost attractiveness of smaller vehicles is counterbalanced by the need for more frequent dump visits, and vice versa. The possibility of open tours, in particular such with different origin and destination depots in a multi-depot setting, is a characteristic that appears less frequently in practice and as a consequence in the literature. Yet, it can lead to important financial savings. Some of the waste collectors in our case study regions need this flexibility but are unable to assess its benefits. This chapter will therefore highlight and quantify the value of such strategies.

2.2 Exact Approach

The formulation we propose introduces several extensions to the model of Sahoo et al. (2005), including multiple origins and destinations, multiple capacities, accessibility restrictions, a maximum tour duration, a richer objective function capturing the costs faced by a real firm, and the elimination of the constraints calculating the necessary number of disposal trips for each vehicle. Unlike in the case of Buhrkal et al. (2012), the driver break is contingent on the start of the tour and not restricted to a time window. In what follows, Section 2.2.1 presents the MILP formulation, while Section 2.2.2 develops several problem-specific valid inequalities shown to have a significant impact on computation time.

2.2.1 Formulation

Formally, we define the problem on a directed graph $\mathcal{G}(\mathcal{N}, \mathcal{A})$, with $\mathcal{N} = \mathcal{O}' \cup \mathcal{O}'' \cup \mathcal{P} \cup \mathcal{D}$, where \mathcal{O}' is the set of origin depots, \mathcal{O}'' is the set of destination depots, \mathcal{P} is the set of containers, \mathcal{D} is the set of dumps, and $\mathcal{A} = \{(i, j) \mid \forall i, j \in \mathcal{N}\}$ is the set of arcs. For modeling purposes, it is assumed that the set \mathcal{D} contains a sufficient number of replications of each dump to allow

multiple visits by the same vehicle. In the computational experiments of Section 2.4.1, we discuss ways of limiting the necessary number of such replications to prevent an explosion of variables.

The set of arcs \mathcal{A} is associated with an asymmetric distance matrix, with π_{ij} the length of arc (i, j) . Each vehicle may have a different average speed, which results in a vehicle-specific travel time matrix, where τ_{ijk} is the travel time of vehicle k on arc (i, j) . Each point has a single time window $[\lambda_i, \mu_i]$, where λ_i and μ_i stand for the earliest and latest possible start-of-service time. Start of service after μ_i is not allowed and if the vehicle arrives before λ_i , it has to wait. Service duration for each point is denoted by δ_i , and the pickup volume and weight by ρ_i^v and ρ_i^w , respectively. The service duration for containers is mostly influenced by the type of container, e.g. underground or overground, and at dumps by factors such as weighing and billing, hence it is not indexed by vehicle. Service duration at the depots is zero.

There is a heterogeneous fixed fleet \mathcal{K} , with each vehicle defined by its capacity in terms of maximum volume Ω_k^v and weight Ω_k^w , a deployment cost φ_k , a unit-distance running cost β_k , and a unit-time running cost θ_k . The vehicle-specific sets $\mathcal{O}'_k \subseteq \mathcal{O}'$ and $\mathcal{O}''_k \subseteq \mathcal{O}''$ designate the available origins and destinations for vehicle k . The set \mathcal{O}'_k degenerates to one point, the current depot, or coincides with \mathcal{O}' in case we need to optimize the vehicle's home depot. Similarly, \mathcal{O}''_k either contains only the depot to which the vehicle is required to return at the end of the day, or many depots to allow for an open tour. Here, we need to stress the difference between a vehicle's home depot and origin depot. The home depot is where the vehicle belongs, while the origin depot is where the vehicle starts a tour, which may be different from the home depot. There is a maximum tour duration H , and a break of duration v must be taken after Υ hours of continuous work, which divides the working day into two roughly equal halves. Accessibility restrictions are described by a binary flag α_{ijk} whose value is 1 if arc (i, j) is accessible for vehicle k , and 0 otherwise.

We introduce the following binary decision variables: $x_{ijk} = 1$ if vehicle k traverses arc (i, j) , 0 otherwise; $r_{ijk} = 1$ if i and j are, respectively, the origin and destination depot of vehicle k , 0 otherwise; $b_{ijk} = 1$ if vehicle k takes a break on arc (i, j) , 0 otherwise; $z_k = 1$ if vehicle k is used, 0 otherwise. Three groups of continuous variables, Q_{ik}^v , Q_{ik}^w and S_{ik} , are defined to track the cumulative volume and weight, and the start-of-service time at point i for vehicle k . Table 2.1 is a summary of the used notations.

The objective function (2.1) minimizes two terms. The first one is the sum of deployment, unit-distance and unit-time running costs for all used vehicles. The second one captures the cost of relocation for open tours. Clearly, open tours lead to reduced cost compared to closed tours since they provide more flexibility. However, the single-day problem ignores the fact that the vehicles will need to return to their home depots at some point in the future. Thus, the goal of the relocation term is to integrate this cost effect, which is multiplied by the weight factor ψ reflecting our degree of conservatism. In effect, the relocation cost term

Table 2.1: Notations

Sets			
\mathcal{O}'	set of origins	\mathcal{O}''	set of destinations
\mathcal{O}'_k	set of origins for vehicle k	\mathcal{O}''_k	set of destinations for vehicle k
\mathcal{P}	set of containers	\mathcal{D}	set of dumps
\mathcal{N}	$= \mathcal{O}' \cup \mathcal{O}'' \cup \mathcal{P} \cup \mathcal{D}$	\mathcal{K}	set of vehicles
Parameters			
φ_k	deployment cost of vehicle k (monetary)		
β_k	unit-distance running cost of vehicle k (monetary)		
θ_k	unit-time running cost of vehicle k (monetary)		
Ω_k^v, Ω_k^w	volume and weight capacity of vehicle k		
π_{ij}	length of arc (i, j)		
τ_{ijk}	travel time of vehicle k on arc (i, j)		
λ_i, μ_i	lower and upper time window bound at point i		
α_{ijk}	1 if arc (i, j) is accessible for vehicle k , 0 otherwise		
δ_i	service duration at point i		
ρ_i^v, ρ_i^w	pickup volume and weight at point i		
H	maximum tour duration		
Υ	maximum continuous work limit after which a break is due		
v	break duration		
ψ	weight of relocation cost term $\in [0, 1]$		
Decision Variables			
x_{ijk}	1 if vehicle k traverses arc (i, j) , 0 otherwise (binary)		
r_{ijk}	1 if i and j are the origin and destination depot of vehicle k , 0 otherwise (binary)		
b_{ijk}	1 if vehicle k takes a break on arc (i, j) , 0 otherwise (binary)		
z_k	1 if vehicle k is used, 0 otherwise (binary)		
Q_{ik}^v	cumulative volume on vehicle k at point i (continuous)		
Q_{ik}^w	cumulative weight on vehicle k at point i (continuous)		
S_{ik}	start-of-service time of vehicle k at point i (continuous)		

incentivizes, rather than enforcing, vehicles to return to their home depots. The motivation behind this comes from our case study, in which there are collectors in wide and sparsely populated rural regions who practice open tours.

$$\min z = \sum_{k \in \mathcal{K}} \left(\varphi_k z_k + \beta_k \sum_{i \in \mathcal{N}} \sum_{j \in \mathcal{N}} \pi_{ij} x_{ijk} + \theta_k \left(\sum_{j \in \mathcal{O}''_k} S_{jk} - \sum_{i \in \mathcal{O}'_k} S_{ik} \right) \right) \quad (2.1)$$

$$+ \psi \sum_{k \in \mathcal{K}} \sum_{i \in \mathcal{O}'_k} \sum_{j \in \mathcal{O}''_k} (\beta_k \pi_{ji} + \theta_k \tau_{jik}) r_{ijk}$$

The constraints can be split into several categories with the first category consisting of basic vehicle routing constraints. Constraints (2.2) impose that each container should be served by exactly one vehicle. Constraints (2.3) and (2.4) ensure that if a vehicle is used, its tour starts at an available origin and ends at an available destination with a visit to a dump immediately before that. Constraints (2.5) forbid entering origins and unavailable destinations and constraints (2.6) forbid leaving destinations and unavailable origins. Accessibility restrictions are enforced by constraints (2.7), while flow conservation is represented by constraints (2.8).

The link between the variables x_{ijk} and r_{ijk} is achieved through constraints (2.9).

$$\sum_{k \in \mathcal{K}} \sum_{j \in \mathcal{P} \cup \mathcal{D}} x_{ijk} = 1, \quad \forall i \in \mathcal{P} \quad (2.2)$$

$$\sum_{i \in \mathcal{O}'_k} \sum_{j \in \mathcal{N}} x_{ijk} = z_k, \quad \forall k \in \mathcal{K} \quad (2.3)$$

$$\sum_{i \in \mathcal{D}} \sum_{j \in \mathcal{O}''_k} x_{ijk} = z_k, \quad \forall k \in \mathcal{K} \quad (2.4)$$

$$\sum_{i \in \mathcal{N}} x_{ijk} = 0, \quad \forall k \in \mathcal{K}, j \in \mathcal{O}' \cup (\mathcal{O}'' \setminus \mathcal{O}''_k) \quad (2.5)$$

$$\sum_{j \in \mathcal{N}} x_{ijk} = 0, \quad \forall k \in \mathcal{K}, i \in \mathcal{O}'' \cup (\mathcal{O}' \setminus \mathcal{O}'_k) \quad (2.6)$$

$$x_{ijk} \leq \alpha_{ijk}, \quad \forall k \in \mathcal{K}, i \in \mathcal{O}'_k \cup \mathcal{P} \cup \mathcal{D}, j \in \mathcal{P} \cup \mathcal{D} \cup \mathcal{O}''_k \quad (2.7)$$

$$\sum_{i \in \mathcal{N}: i \neq j} x_{ijk} = \sum_{i \in \mathcal{N}: i \neq j} x_{jik}, \quad \forall k \in \mathcal{K}, j \in \mathcal{P} \cup \mathcal{D} \quad (2.8)$$

$$\sum_{m \in \mathcal{N}} x_{imk} + \sum_{m \in \mathcal{D}} x_{mjk} - 1 \leq r_{ijk}, \quad \forall k \in \mathcal{K}, i \in \mathcal{O}'_k, j \in \mathcal{O}''_k \quad (2.9)$$

In the context of vehicle capacities, constraints (2.10) and (2.11) limit, respectively, the cumulative volume and weight on the vehicle at each point, while constraints (2.12) and (2.13) reset them to zero at the dumps, origins and destinations. Keeping track of the cumulative volume and weight on the vehicle is achieved by constraints (2.14) and (2.15).

$$\rho_i^v \leq Q_{ik}^v \leq \Omega_k^v, \quad \forall k \in \mathcal{K}, i \in \mathcal{P} \quad (2.10)$$

$$\rho_i^w \leq Q_{ik}^w \leq \Omega_k^w, \quad \forall k \in \mathcal{K}, i \in \mathcal{P} \quad (2.11)$$

$$Q_{ik}^v = 0, \quad \forall k \in \mathcal{K}, i \in \mathcal{N} \setminus \mathcal{P} \quad (2.12)$$

$$Q_{ik}^w = 0, \quad \forall k \in \mathcal{K}, i \in \mathcal{N} \setminus \mathcal{P} \quad (2.13)$$

$$Q_{ik}^v + \rho_j^v \leq Q_{jk}^v + \Omega_k^v (1 - x_{ijk}), \quad \forall k \in \mathcal{K}, i \in \mathcal{O}'_k \cup \mathcal{P} \cup \mathcal{D}, j \in \mathcal{P} \quad (2.14)$$

$$Q_{ik}^w + \rho_j^w \leq Q_{jk}^w + \Omega_k^w (1 - x_{ijk}), \quad \forall k \in \mathcal{K}, i \in \mathcal{O}'_k \cup \mathcal{P} \cup \mathcal{D}, j \in \mathcal{P} \quad (2.15)$$

The next four constraints express the temporal characteristics of the problem. Constraints (2.16) calculate the start-of-service time at each point, including service duration and a possible break duration. In addition, these constraints eliminate the possibility of subtours and ensure that a point will not be visited more than once by the same vehicle. The value of M_1 can be set to $\mu_i + \delta_i + v + \tau_{ijk}$. Constraints (2.17), (2.18) and (2.19) enforce the time windows and maximum tour duration. We assume that the lower time window bound is restrictive at the origins and the upper one at the destinations.

$$S_{ik} + \delta_i + v b_{ijk} + \tau_{ijk} \leq S_{jk} + M_1 (1 - x_{ijk}), \quad \forall k \in \mathcal{K}, i \in \mathcal{O}'_k \cup \mathcal{P} \cup \mathcal{D}, j \in \mathcal{P} \cup \mathcal{D} \cup \mathcal{O}''_k \quad (2.16)$$

$$\lambda_i \sum_{j \in \mathcal{N}} x_{ijk} \leq S_{ik}, \quad \forall k \in \mathcal{K}, i \in \mathcal{O}'_k \cup \mathcal{P} \cup \mathcal{D} \quad (2.17)$$

$$S_{jk} \leq \mu_j \sum_{i \in \mathcal{N}} x_{ijk}, \quad \forall k \in \mathcal{K}, j \in \mathcal{P} \cup \mathcal{D} \cup \mathcal{O}''_k \quad (2.18)$$

$$0 \leq \sum_{j \in \mathcal{O}''_k} S_{jk} - \sum_{i \in \mathcal{O}'_k} S_{ik} \leq H, \quad \forall k \in \mathcal{K} \quad (2.19)$$

The next block of constraints determines the arc on which a break is due. Breaks are modeled on the arcs as in much of the vehicle routing literature and can in practice be taken on the arcs' tails. Constraints (2.20) and (2.21) limit the arcs on which the break can be taken so that it is taken as late as possible. The value of M_2 can be fixed as $\max(\mu_i - \min_{m \in \mathcal{O}'_k} \lambda_m + \delta_i - \Upsilon, 0)$ and that of M_3 as $\Upsilon + \max_{m \in \mathcal{O}'_k} \mu_m$. Constraints (2.22) impose that the vehicle can only take a break on the arcs it traverses. Finally, constraints (2.23) ensure that the break is actually taken if the vehicle tour is longer than the maximum continuous work limit Υ , after which a break is due.

$$\left(S_{ik} - \sum_{m \in \mathcal{O}'_k} S_{mk} \right) + \delta_i - \Upsilon \leq M_2(1 - b_{ijk}), \quad \forall k \in \mathcal{K}, i \in \mathcal{O}'_k \cup \mathcal{P} \cup \mathcal{D}, j \in \mathcal{P} \cup \mathcal{D} \cup \mathcal{O}''_k \quad (2.20)$$

$$\Upsilon - \left(S_{jk} - \sum_{m \in \mathcal{O}'_k} S_{mk} \right) \leq M_3(1 - b_{ijk}), \quad \forall k \in \mathcal{K}, i \in \mathcal{O}'_k \cup \mathcal{P} \cup \mathcal{D}, j \in \mathcal{P} \cup \mathcal{D} \cup \mathcal{O}''_k \quad (2.21)$$

$$b_{ijk} \leq x_{ijk}, \quad \forall k \in \mathcal{K}, i, j \in \mathcal{N} \quad (2.22)$$

$$\left(\sum_{j \in \mathcal{O}''_k} S_{jk} - \sum_{i \in \mathcal{O}'_k} S_{ik} \right) - \Upsilon \leq (H - \Upsilon) \sum_{i \in \mathcal{N}} \sum_{j \in \mathcal{N}} b_{ijk}, \quad \forall k \in \mathcal{K} \quad (2.23)$$

Finally, constraints (2.24) to (2.26) establish the variable domains.

$$x_{ijk}, b_{ijk}, z_k \in \{0, 1\}, \quad \forall k \in \mathcal{K}, i, j \in \mathcal{N} \quad (2.24)$$

$$r_{ijk} \in \{0, 1\}, \quad \forall k \in \mathcal{K}, i \in \mathcal{O}', j \in \mathcal{O}'' \quad (2.25)$$

$$Q_{ik}^v, Q_{ik}^w, S_{ik} \geq 0, \quad \forall k \in \mathcal{K}, i \in \mathcal{N} \quad (2.26)$$

2.2.2 Variable Fixing and Valid Inequalities

We can exploit the special structure of our problem by fixing some of the binary variables and defining several valid inequalities that restrict the search space of some of the binary and continuous variables without eliminating any feasible solutions. We first set to zero binary variables linked to impossible traversals. Constraints (2.27) eliminate the possibility of loops. In a similar fashion, constraints (2.28), (2.29) and (2.30) forbid traveling from an origin to a dump or destination, from a container to a destination, and from a dump to another dump,

respectively.

$$x_{iik} = 0, \quad \forall k \in \mathcal{K}, i \in \mathcal{N} \quad (2.27)$$

$$x_{ijk} = 0, \quad \forall k \in \mathcal{K}, i \in \mathcal{O}'_k, j \in \mathcal{D} \cup \mathcal{O}''_k \quad (2.28)$$

$$x_{ijk} = 0, \quad \forall k \in \mathcal{K}, i \in \mathcal{P}, j \in \mathcal{O}''_k \quad (2.29)$$

$$x_{ijk} = 0, \quad \forall k \in \mathcal{K}, i \in \mathcal{D}, j \in \mathcal{D}: i \neq j \quad (2.30)$$

The presence of time windows allows us to fix time-window infeasible traversals. Constraints (2.31) express the fact that if by visiting point i as early as possible vehicle k cannot visit point j within its time window, then points i and j cannot be visited by the same vehicle k , i.e. arc (i, j) is not traversed by vehicle k . These first two sets of rules can also be used to eliminate all the big M constraints (2.14, 2.15, 2.16, 2.20, 2.21) for such variables as they become trivial.

$$x_{ijk} = 0, \quad \forall k \in \mathcal{K}, i \in \mathcal{O}'_k \cup \mathcal{P} \cup \mathcal{D}, j \in \mathcal{P} \cup \mathcal{D} \cup \mathcal{O}''_k: \lambda_i + \delta_i + \tau_{ijk} > \mu_j \quad (2.31)$$

The first set of valid inequalities is used to restrict the start-of-service time search space. Inequalities (2.32) impose a lower bound, short of waiting times, on the difference between the start-of-service time at the origin and destination for each used vehicle. Then inequalities (2.33) and (2.34) calculate the latest possible start and earliest possible finish of each tour. In constraints (2.33) $\mathcal{P}^a \subseteq \mathcal{P}$, s.t. $\alpha_{imk} = 1$, and in constraints (2.34) $\mathcal{D}^a \subseteq \mathcal{D}$, s.t. $\alpha_{mjk} = 1$.

$$\sum_{j \in \mathcal{O}''_k} S_{jk} - \sum_{i \in \mathcal{O}'_k} S_{ik} \geq \sum_{i \in \mathcal{N}} \sum_{j \in \mathcal{N}} x_{ijk} (\delta_i + \tau_{ijk}), \quad \forall k \in \mathcal{K} \quad (2.32)$$

$$S_{ik} \leq \max_{m \in \mathcal{P}^a} (\mu_m - \tau_{imk}) z_k, \quad \forall k \in \mathcal{K}, i \in \mathcal{O}'_k \quad (2.33)$$

$$S_{jk} \geq \min_{m \in \mathcal{D}^a} (\lambda_m + \delta_m + \tau_{mjk}) \sum_{m \in \mathcal{D}^a} x_{mjk}, \quad \forall k \in \mathcal{K}, j \in \mathcal{O}''_k \quad (2.34)$$

If the problem involves subsets of identical vehicles, the presence of symmetry can substantially reduce the effectiveness of the model. Let $\mathcal{K}' \subseteq \mathcal{K}$ represent a subset of identical vehicles and let $k'_g \in \mathcal{K}'$, where $g \in 1, \dots, |\mathcal{K}'|$ introduces a simple ordering of the elements of \mathcal{K}' . Then for each subset \mathcal{K}' we apply constraints (2.35) or (2.36). These symmetry-breaking constraints specify that the first vehicle in \mathcal{K}' executes the tour with the highest waste volume (weight), the second vehicle executes the tour with the second highest waste volume (weight), etc.

$$\sum_{i \in \mathcal{P}} \sum_{j \in \mathcal{P} \cup \mathcal{D}} \rho_i^v x_{ijk'_g} \geq \sum_{i \in \mathcal{P}} \sum_{j \in \mathcal{P} \cup \mathcal{D}} \rho_i^v x_{ijk'_{g+1}}, \quad \forall g \in 1, \dots, (|\mathcal{K}'| - 1) \quad (2.35)$$

$$\sum_{i \in \mathcal{P}} \sum_{j \in \mathcal{P} \cup \mathcal{D}} \rho_i^w x_{ijk'_g} \geq \sum_{i \in \mathcal{P}} \sum_{j \in \mathcal{P} \cup \mathcal{D}} \rho_i^w x_{ijk'_{g+1}}, \quad \forall g \in 1, \dots, (|\mathcal{K}'| - 1) \quad (2.36)$$

Symmetries will also result from the fact that dumps are replicated to allow multiple visits of the same dump by a particular vehicle. Here, we define $\mathcal{D}' \subseteq \mathcal{D}$ as a subset of replications of

the same physical dump and let $j'_g \in \mathcal{D}'$, where $g \in 1, \dots, |\mathcal{D}'|$ introduces a simple ordering of the elements of \mathcal{D}' . Then, for each subset \mathcal{D}' we apply the lexicographical ordering constraints (2.37), which stipulate that a dump with a higher index should be preceded by a container with a higher index.

$$\sum_{i \in \mathcal{P}} i x_{i j'_g k} \leq \sum_{i \in \mathcal{P}} i x_{i j'_{g+1} k}, \quad \forall k \in \mathcal{K}, g \in 1, \dots, (|\mathcal{D}'| - 1) \quad (2.37)$$

The last set of valid inequalities concerns the dump visits. With (2.38) we impose that a dump replication may be visited at most once by a vehicle. With (2.39), on the other hand, we set for every vehicle the maximum number of trips from dumps to containers. It takes into account the fact that the first trip to a container is from the origin depot and the last trip from a dump is to the destination depot.

$$\sum_{i \in \mathcal{P}} x_{i j k} \leq 1, \quad \forall k \in \mathcal{K}, j \in \mathcal{D} \quad (2.38)$$

$$\sum_{i \in \mathcal{D}} \sum_{j \in \mathcal{P}} x_{i j k} \leq \min(|\mathcal{D}| - 1, |\mathcal{P}| - 1), \quad \forall k \in \mathcal{K} \quad (2.39)$$

With the addition of the valid inequalities, a state-of-the-art MIP solver like Gurobi can handle instances with 10-15 containers, a depot, two to eight dumps before replication, and six vehicles, with the only critical resource being computation time. Computation times are influenced both by the instance sizes and by their spatial and temporal characteristics. We return to this question in Section 2.4.1.

2.3 Meta-heuristic Approach

The vehicle routing problem is well known to be *NP*-hard (see e.g. Garey and Johnson, 1979). Being a generalization thereof, our waste collection problem is even harder to solve. Moreover, realistic instances involving 50 or more containers and several depots, dumps and vehicles will translate into thousands of binary variables and tens of thousands of constraints. Therefore, we develop a meta-heuristic approach based on Multiple Neighborhood Search (MNS), which is capable of systematically treating all problem features. Section 2.3.1 below defines solution feasibility. This is followed by the description of the initial construction procedure in Section 2.3.2 and the iterative MNS algorithm in Section 2.3.3.

2.3.1 Feasibility

A solution to our problem is a set of tours in which breaks have been inserted. It is considered feasible if all tours that comprise it satisfy four criteria. First, start-of-service times should respect time windows. Secondly, tour duration should be shorter than or equal to the maximum tour duration. These two criteria may be thought of as expressing temporal feasibility. Thirdly, the volume and weight capacities of the vehicles may not be violated at

any point. This can be ensured by inserting appropriate visits to the available dumps. We attach to this last criterion the condition that a tour should start and finish at an available depot and visit a dump before the end. Finally, accessibility restrictions should be respected.

Every insertion or removal of a point from a tour, and every application of a neighborhood operator requires the recalculation of start-of-service and waiting times for all or part of the points in the tour. As shown in Algorithm 2.1, we consider a tour served by vehicle $k \in \mathcal{K}$, for brevity tour k , represented as an ordered sequence of points $1, 2, \dots, n-1, n$ indexed by i . The calculation begins by setting the start-of-service time at the origin, S_{1k} , as early

Algorithm 2.1: Temporal Feasibility Algorithm

Input tour k as a sequence of points $1, 2, \dots, n-1, n$

Output start-of-service times, waiting times and temporal feasibility of tour k

```

1:  $S_{1k} \leftarrow \lambda_1$ 
2: for  $i = 2, 3, \dots, n-1, n$  in tour  $k$  do
3:    $S_{ik} \leftarrow S_{(i-1)k} + \delta_{i-1} + \tau_{(i-1)ik}$ 
4:   if  $S_{(i-1)k} < S_{1k} + \Upsilon$  and  $S_{ik} + \delta_i > S_{1k} + \Upsilon$  then
5:      $S_{ik} \leftarrow S_{ik} + \upsilon$ 
6:   end if
7:   if  $S_{ik} < \lambda_i$  then
8:      $w_{ik} \leftarrow \lambda_i - S_{ik}$ 
9:      $S_{ik} \leftarrow \lambda_i$ 
10:  else
11:     $w_{ik} \leftarrow 0$ 
12:  end if
13: end for
14: if  $S_{ik} \leq \mu_i, \forall i$  then
15:   for  $i = n, n-1, \dots, 3, 2$  in tour  $k$  do
16:    if  $w_{ik} > 0$  then
17:       $S'_{(i-1)k} \leftarrow S_{(i-1)k}$ 
18:       $S_{(i-1)k} \leftarrow \min(S_{(i-1)k} + w_{ik}, \mu_{i-1})$ 
19:       $w_{(i-1)k} \leftarrow w_{(i-1)k} + (S_{(i-1)k} - S'_{(i-1)k})$ 
20:       $w_{ik} \leftarrow w_{ik} - (S_{(i-1)k} - S'_{(i-1)k})$ 
21:    end if
22:  end for
23:   $w_{1k} \leftarrow 0$ 
24:  if  $S_{nk} - S_{1k} \leq H$  then
25:    tour  $k$  is temporally feasible
26:  else
27:    tour  $k$  is duration infeasible
28:  end if
29: else
30:   tour  $k$  is time-window infeasible
31: end if

```

as possible. For each subsequent point i , S_{ik} is tentatively calculated as the sum of the start-of-service time at point $i-1$, the service duration at point $i-1$, and the travel time from $i-1$ to i , i.e. $S_{ik} = S_{(i-1)k} + \delta_{i-1} + \tau_{(i-1)ik}$. If the maximum continuous working time limit Υ expires between the start-of-service time at $i-1$ and the end-of-service time at i , in other words if $S_{(i-1)k} < S_{1k} + \Upsilon$ and $S_{ik} + \delta_i > S_{1k} + \Upsilon$, we need to insert the required break before serving point i , which is achieved by incrementing S_{ik} by the break duration v . Finally, if S_{ik} violates the lower time window bound λ_i , i.e. if $S_{ik} < \lambda_i$, we introduce waiting time w_{ik} at point i , equal to the difference $\lambda_i - S_{ik}$, and update S_{ik} to λ_i . Once all S_{ik} have been determined, we check if upper time window bounds μ_i are respected for all i . If this is the case, we apply forward time slack reduction on the tour, otherwise we declare the tour time-window infeasible.

Forward time slack, as described by Savelsbergh (1992), keeps track of the maximum amount each start-of-service time can be delayed without violating time windows on the tour. We examine points sequentially in reverse order. If there is waiting at point i , there could be a non-zero slack at point $i-1$, because pushing $S_{(i-1)k}$ forward may eliminate or reduce waiting at i . We can push $S_{(i-1)k}$ forward by the amount of waiting at i , or until we reach the upper time window bound at $i-1$. The last operation is expressed as $S_{(i-1)k} = \min(S_{(i-1)k} + w_{ik}, \mu_{i-1})$, and it entails an update of $w_{(i-1)k}$ and w_{ik} to factor in the potential increase of waiting at $i-1$ and decrease of waiting at i . Let $S'_{(i-1)k}$ denote the original start-of-service time at point $i-1$ before slack reduction. Then, waiting at $i-1$ will be increased by the difference between $S_{(i-1)k}$ and $S'_{(i-1)k}$, and waiting at i will be reduced by the same difference. Finally, we need to artificially put $w_{1k} = 0$. Forward time slack reduction preserves time-window feasibility. Therefore, after the procedure it only remains to check if the tour's duration is feasible. If it is the case, we accept the tour as temporally feasible, otherwise we declare it duration infeasible.

Verifying capacity feasibility is much more straightforward. At each point of the tour, we calculate the cumulative volume and weight loads, Q_{ik}^v and Q_{ik}^w , on the vehicle, resetting both to zero if the point is a dump. If, for any point i , $Q_{ik}^v > \Omega_k^v$ or $Q_{ik}^w > \Omega_k^w$ or a dump is not visited immediately before the destination, we declare the tour capacity infeasible. The logic behind accessibility feasibility is trivial. Implementation-wise, we construct vehicle tours only using accessible points. Inaccessibilities may occur with the application of inter-tour operators and a simple count of the number of inaccessible points is updated with the application of an operator. The latter is much more efficient than inspecting the points when accessibility feasibility needs to be verified.

2.3.2 Initial Solution Construction

Tour construction is performed sequentially. Initially, all containers belong to the pool of unassigned containers \mathcal{P} , and all vehicles to the pool of unassigned vehicles \mathcal{K} . A seed tour is created by assigning the cheapest feasible sequence of origin, container, dump and

destination to the cheapest available vehicle. This assumes that accessible containers remain for the current vehicle; otherwise unassigned containers are swapped with containers already in other tours until an accessible container is found that can be feasibly inserted in the current tour. All assigned vehicles and containers are removed from their respective unassigned pools. Once a seed tour k has been created, it is expanded using a feasibility preserving best insertion heuristic. At each iteration, we insert container $i \in \mathcal{P}$ at the position j in the tour that would yield the smallest cost increase. The point at position j , as well as all subsequent points, are shifted to the right.

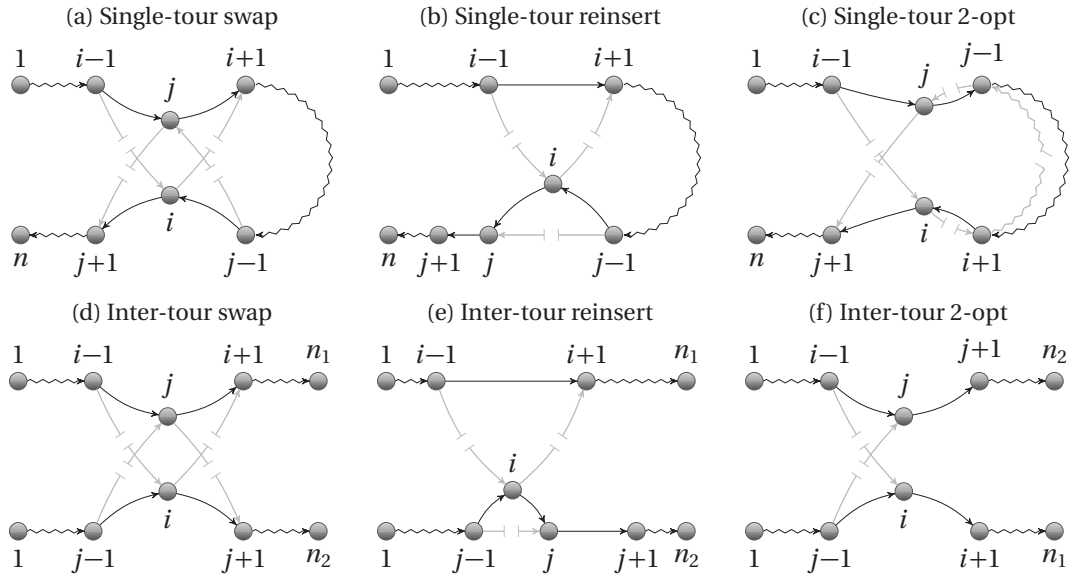
If no more feasible container insertions are possible and if infeasibility would result from capacity violation, we insert a dump using the same feasibility preserving best insertion logic, otherwise we terminate the tour. In the former case, we insert dump $i \in \mathcal{D}$ at the position j in the tour that leads to the smallest cost increase. In addition, we require that the dump not be inserted as an immediate predecessor or successor of another dump on the tour or just after the origin depot. If there are no feasible dump insertions, the tour is terminated. Finally, to avoid a meaningless increase in the objective function, we require that after a dump insertion there should be at least one feasible container insertion. If this condition does not hold, the last inserted dump is removed and the tour is terminated.

Tour construction stops when the pool of unassigned containers is empty, or the pool of unassigned vehicles is empty, or infeasibilities prohibit further insertions. When each tour is constructed, it is individually improved using the single-tour operators described in Section 2.3.3 below. There is no guarantee that the construction procedure will be able to insert all containers. Therefore, additional insertions are periodically attempted in the MNS, as described in Section 2.3.3. Thus, if the underlying problem is feasible, as are all the benchmarks we use, and the construction procedure has not inserted all containers, they should be able to be inserted by the MNS. If, however, the underlying problem is infeasible, as in a realistic application where an insufficient fleet is provided for a large number of containers, the MNS will still produce a solution respecting all constraints for the inserted points, even if some containers remain unserved.

2.3.3 Multiple Neighborhood Search

In order to keep the improvement phase as general as possible, we consider three neighborhoods—swap, reinsert, and 2-opt, with each neighborhood using classical single- and inter-tour operators of the respective type. Figure 2.1 depicts the six operators with possible improvements from the application of each of them. The interrupted gray arcs form parts of the tours before the application of the operators. The resulting improved tours are given in solid black arcs. The application of an operator, whether single- or inter-tour, may lead to a feasible or an infeasible neighbor. If the neighbor solution is infeasible, its objective function (2.1) is multiplied by a factor *infFactor* larger than one. If the next neighbor is feasible, this factor is reduced by *infStepDown*, and if infeasible, it is increased by *infStepUp*. The factor

Figure 2.1: Multiple Neighborhood Search Operators



Note. This figure depicts improvements that can result from the application of each operator, with the interrupted gray arcs replaced by the solid black arcs.

infFactor will never drop below one.

To define the operators more precisely, the single-tour swap disconnects $i-1$ from i , i from $i+1$, $j-1$ from j , and j from $j+1$, and reconnects $i-1$ to j , j to $i+1$, $j-1$ to i , and i to $j+1$. Its inter-tour version works in exactly the same way with the only difference being that i and j belong to different tours. The single-tour reinsert disconnects $i-1$ from i , and i from $i+1$, reconnecting $i-1$ directly to $i+1$. Then it disconnects $j-1$ from j and reconnects $j-1$ to i , and i to j . The logic of the inter-tour reinsert is the same, with i and j belonging to different tours. In essence, the last two operators remove a point i from its original position and insert it in the position of another point j , from the same or a different tour, pushing j to the right. Finally, the single-tour 2-opt disconnects $i-1$ from i , and j from $j+1$, and reconnects $i-1$ to j , and i to $j+1$, thus reversing the orientation of the section $i, i+1, \dots, j-1, j$, inclusive of i and j . The inter-tour 2-opt disconnects $i-1$ from i , and $j-1$ from j , where i and j belong to different tours, and reconnects $i-1$ to j , and $j-1$ to i , which results in the exchange of the end portions of the two affected tours, inclusive of i and j .

As described in Algorithm 2.2, the succession of neighborhoods (swap, reinsert, 2-opt in that order) is applied until either *maxIter* iterations or *maxNonImplIter* non-improving iterations has been reached. Each individual neighborhood is applied for *maxNbIter* iterations or *maxNbNonImplIter* non-improving iterations from the last visited local minimum, and at each neighborhood change we start again from the best feasible solution found so far and reset *infFactor*. For each neighbor, a random sample of single- and inter-tour moves of the current neighborhood is evaluated and the cheapest one, feasible or infeasible, in the

last case evaluated after multiplication by $infFactor$, is accepted as the new incumbent. To prevent cycling and encourage diversification towards less explored areas of the search space, a solution with the same objective value is not admitted more than once for a given number of iterations, denoted by $cycleFreq$. These non-admissible solutions are held in a ban list. When a new incumbent is generated, the ban list is updated to include its cost and exclude the cost older than $cycleFreq$.

In each neighborhood, at each $recoverFreq$ iterations, and if the available fleet is heterogeneous (i.e. at least one vehicle is different from the rest), we evaluate and perform vehicle

Algorithm 2.2: Multiple Neighborhood Search

Define: \mathcal{K} is the set of all available vehicles

Input set of constructed tours $\mathcal{K}' \subseteq \mathcal{K}$

Output set of improved tours $\mathcal{K}'' \subseteq \mathcal{K}$

```

1: initialize  $infFactor$ 
2: initialize ban list
3: initialize start neighborhood
4:  $currentIncumbent \leftarrow$  solution from tour construction
5: for  $maxIter$  do
6:   for each neighborhood do
7:     for  $maxNbIter$  do
8:        $\mathcal{N} \leftarrow$  random neighbor sample of  $currentIncumbent$ 
9:        $currentIncumbent \leftarrow \min(n)\{\text{cost}(n) \mid \forall n \in \mathcal{N} : \text{cost}(n) \notin \text{ban list}\}$ 
10:      update  $infFactor$ 
11:      update ban list
12:      if reached  $recoverFreq$  then
13:        vehicle reasg eval procedure with cap recovery and depot reasg.
14:        improve tours individually
15:        update ban list
16:      end if
17:      if reached  $maxNbNonImplter$  then
18:        change neighborhood
19:        reset  $currentIncumbent$  to best feasible solution found so far
20:        reset  $infFactor$ 
21:        break
22:      end if
23:    end for
24:    change neighborhood
25:    reset  $currentIncumbent$  to best feasible solution found so far
26:    reset  $infFactor$ 
27:  end for
28:  if reached  $maxNonImplter$  then
29:    break
30:  end if
31: end for

```

reassignments to tours. The vehicle reassignment evaluation procedure unassigns all assigned vehicles and starts inspecting the tours in a descending order of total load. For each consecutively examined tour, the best vehicle is assigned so that the assignment is feasible. The assignment feasibility is verified after applying the capacity recovery procedure described next. If no feasible assignment is possible, the available vehicle with the largest capacity is assigned to the tour. After the assignment, the capacity recovery procedure is rerun and the tour is individually improved. Alternatively, if the available fleet is homogeneous, we proceed directly to the capacity recovery procedure, followed by individual improvement of all tours.

The logic of the vehicle reassignment evaluation procedure is somewhat different compared to what would benefit Taillard's (1999) HFFVRP formulation. Here the procedure tries to balance between two conflicting goals. Assuming a fleet with correlated characteristics, assigning cheaper vehicles to tours is counterbalanced by the necessity for more frequent visits to the dumps, because cheaper vehicles have smaller capacities. This is compounded by the fact that, in a realistic scenario, dumps (intermediate facilities) are located outside the collection area (for example in suburbs or industrial zones) instead of centrally as in the benchmark instances we see. Therefore, the logic of this procedure is different, as is its direct applicability to a pure HFFVRP formulation where reassignments have to be examined with every move. In our case, capacity infeasibility resulting after a move may be recovered by adding more dump visits or simply reordering them. Moreover, if a tour is attracting points from other tours, reassigning vehicles too often may have an adverse effect.

The purpose of the capacity recovery procedure is the restoration of capacity feasibility through the reordering of dump visits. Such an action would be necessary, for example, when a new vehicle with a smaller capacity is assigned to a tour. In a broader sense, it also serves to test dump visits that are not present in the current solution or exclude dump visits that have become redundant. The procedure starts by removing all dump visits from the tour. Then it determines the minimum number of necessary visits by inspecting consecutively the points in the tour and inserting a dump visit at the best position before capacity is exceeded. Thus, it removes unnecessary dump visits from short tours and inserts additional dump visits or reorders the dump visits in tours that may have been rendered capacity infeasible by the neighborhood operators or the assignment of a new vehicle.

This is followed by the reassignment evaluation of destination depots. Given the generally small number of available destination depots, all possibilities are evaluated. The subsequent individual improvement may be able to recover infeasibilities related to maximum tour duration or time window violations. Moreover, if unassigned containers remain and can be feasibly inserted, new insertions are attempted during individual tour improvement before switching back from 2-opt to swap. There is a penalty associated with unassigned containers, which encourages assignment with a near-guarantee of cost improvement. In the end, the logic of this meta-heuristic approach is such that it remains fairly general rather than being tailored to narrowly specified problem instances.

2.4 Numerical Experiments

The MNS algorithm is coded as a single-thread application in Java and the mathematical model is solved using the Gurobi 6.0.0 MIP solver via its Java API. The MNS uses the parameter values presented in Table 2.2, which are selected after extensive trial and adjustment on the instance sets. In the experiments below, each instance is solved 10 times and Gurobi is warm-started with the best solution out of the 10 runs. All tests are performed on a 3.20 GHz Intel Core i5 desktop computer with 8 GB of memory running a 64-bit Windows 7.

Table 2.2: Algorithmic Parameters

Parameter	Value	Parameter	Value
<i>maxNbIter</i>	30	<i>infStepUp</i>	0.05
<i>maxNbNonImpIter</i>	7	<i>infStepDown</i>	0.02
<i>maxIter</i>	100 (10 ^a)	<i>cycleFreq</i>	∞
<i>maxNonImpIter</i>	15 (3 ^a)	<i>recoverFreq</i>	5
<i>infFactor</i>	1.10	sample size	10 ^b

^a *Note.* Value for individual tour improvement.

^b *Note.* At a given iteration, the chosen operator is evaluated on each point i for a random sample of 10 j points, see Figure 2.1.

In the following, Section 2.4.1 compares the MNS to Gurobi on modifications of the small Schneider et al. (2014) E-VRPTW instances, while Section 2.4.2 tests the MNS on the Best Known Solutions (BKS) to Crevier et al.'s (2007) MDVRPI instance sets. On the same sets, we also evaluate the benefit of open tours with different origin and destination depots. The BKS to Crevier et al.'s (2007) MDVRPI instances are due to Hemmelmayr et al. (2013) who use a 2.4 GHz machine with 4 GB of memory, but the processor type is not specified. Moreover, the two algorithms run on different platforms and thus scaling of computation times will almost certainly be biased. Therefore, we report the original computation times with the remark that all results are produced on contemporary processor architectures. Finally, Section 2.4.3 presents a case study of a recyclable waste collector in the canton of Geneva, Switzerland, for which we report significant improvements to the state of practice

2.4.1 Evaluation on Small Instances

For the experiments here, we modified the small Schneider et al. (2014) E-VRPTW instances by adding features that appear in our problem, while ignoring those that are irrelevant. These are 36 instances split into three sets of 12 instances, with five, 10 and 15 customers, respectively. They are derived from the Solomon (1987) instances for the VRP with time windows and thus each set contains instances derived from the R (random customer distribution), C (clustered customer distribution) and RC (mixture of both) classes. Furthermore, subsets R1, C1 and RC1 have a short scheduling horizon, while subsets R2, C2 and RC2 have a long

scheduling horizon, which has an impact on the number of vehicles required to serve all customers. Schneider et al. (2014) add one recharging station at the depot of each instance, as well as additional recharging stations at random locations. The total number of recharging stations is two to eight.

Schneider et al. (2014) also set other parameters related to the vehicle's battery capacity, fuel consumption and inverse refueling rate. However, since these are irrelevant for our problem, they are not discussed further. We assume the same setup of Euclidean distances with speed equal to one. For our purposes, we regard the recharging facilities as dumps and modify the instances by including two vehicle classes. Class A has a capacity of 100, a deployment cost of 50 and a unit distance running cost of one. No time related cost is considered. Class B has a capacity of 120 and its costs are 120% of those of class A. Moreover, in each instance there are two points that are inaccessible for class B vehicles. For each instance, there are three vehicles of each class available, thus imitating a situation of a fixed fleet that handles instances representing different days. Since class B vehicles cannot serve all points, for smaller instances and instances with longer scheduling horizons, class A vehicles will tend to be favored. We also impose a maximum tour duration of 500 and a maximum continuous work limit of 250, after which a break of duration 50 is due.

As briefly discussed in Section 2.2, to solve these instances with the MILP formulation, we need to replicate the dumps a sufficient number of times. There are various approaches to tackling this issue. In general, tighter time windows should be beneficial for computation time. Therefore, one approach is to generate replications with successive time windows of a length that makes it impossible to visit a replication twice. The latter can be calculated, for example, as the minimum travel time from a dump to a container and back to the same dump, including service durations. However, if the scheduling horizon is long compared to the individual travel times among the points, as it is for the R2, C2 and RC2 subsets, the number of such replications becomes excessive. In our experiments, for some instances, these replications can be in the hundreds, which leads to an explosion of variables.

Disregarding successive time windows, the number of necessary dump replications is bounded by the number of containers to serve, which may again be large. Therefore, we use a rule of thumb, where the number of replications for each dump R_d is set as:

$$R_d = \left\lceil \frac{\sum_{i \in \mathcal{P}} \rho_i}{0.75 \Omega_{k^*}} \right\rceil. \quad (2.40)$$

In the above expression, Ω_{k^*} designates the capacity of the smallest vehicle. We omit the superscript v or w because only a single dimension is assumed for the commodity being collected in these instances. This rule states that each dump is replicated a sufficient number of times so that if there is only one tour that is executed by the smallest vehicle, the latter is emptied on average when it is 75% full if it always visits the same dump.

Table 2.3 presents a comparison of the results of the MNS and the results obtained by Gurobi

Table 2.3: MNS vs Solver on Modified Schneider et al. (2014) Instances

Instance	Vehicles	MNS				Solver on Model with Valid Inequalities				Solver on Model Without Valid Inequalities			
		Best	Avg	Runtime Avg (s.)	Objective	MIP Gap (%)	Runtime (s.)	Improvement (%)	Objective	MIP Gap (%)	Runtime (s.)	Improvement (%)	
c101C5	4	489.70	489.70	0.05	489.70	0.00	0.39	0.00	489.70	35.71	7200.01	0.00	
c103C5	2	281.33	281.33	0.04	268.09	0.00	0.17	-4.94	268.09	0.00	4910.40	-4.94	
c206C5	2	374.67	374.67	0.06	360.09	0.00	0.24	-4.05	360.09	0.00	90.27	-4.05	
c208C5	2	343.20	343.20	0.06	343.20	0.00	0.49	0.00	343.20	38.57	7200.04	0.00	
r104C5	1	182.81	182.81	0.02	182.81	0.00	0.04	0.00	182.81	0.00	1.90	0.00	
r105C5	2	251.15	251.15	0.05	251.15	0.00	0.08	0.00	251.15	0.00	0.22	0.00	
r202C5	1	176.52	176.52	0.02	176.52	0.00	0.05	0.00	176.52	0.00	2.67	0.00	
r203C5	1	228.05	228.05	0.03	228.05	0.00	0.12	0.00	228.05	0.00	1504.85	0.00	
rc105C5	2	327.19	327.19	0.05	327.19	0.00	0.15	0.00	327.19	25.27	7200.04	0.00	
rc108C5	2	345.87	345.87	0.04	345.87	0.00	0.15	0.00	345.87	0.00	2069.22	0.00	
rc204C5	1	223.17	223.17	0.09	223.17	0.00	0.17	0.00	223.17	0.00	1327.76	0.00	
rc208C5	1	212.67	212.67	0.02	212.67	0.00	0.25	0.00	212.67	0.00	1156.35	0.00	
c101C10	6	846.10	846.10	0.61	837.13	0.00	5489.48	-1.07	846.10	77.21	7200.48	0.00	
c104C10	3	456.86	456.86	0.43	456.86	0.00	37.28	0.00	456.86	53.05	7200.08	0.00	
c202C10	4	549.74	549.74	0.42	549.74	18.44	7200.09	0.00	549.74	67.71	7200.31	0.00	
c205C10	4	568.92	568.92	0.58	568.58	0.00	2788.37	-0.06	568.92	64.77	7200.11	0.00	
r102C10	3	391.14	391.14	0.40	391.14	0.00	158.70	0.00	391.14	47.69	7200.16	0.00	
r103C10	2	288.67	288.67	0.50	288.67	0.00	18.39	0.00	288.67	43.72	7200.04	0.00	
r201C10	2	310.16	310.16	0.45	310.16	0.00	45.22	0.00	310.16	43.64	7200.46	0.00	
r203C10	2	329.78	329.78	1.13	329.78	0.00	5757.28	0.00	329.78	47.26	7200.08	0.00	
rc102C10	3	534.75	534.75	0.40	534.75	0.00	6.25	0.00	534.75	38.77	7200.09	0.00	
rc108C10	2	429.79	429.79	0.42	429.79	0.00	6.94	0.00	429.79	25.30	7200.09	0.00	
rc201C10	2	502.45	502.45	0.40	498.88	0.00	147.43	-0.51	502.45	58.48	7200.09	0.00	
rc205C10	2	428.80	428.80	0.45	421.36	0.00	26.00	-1.77	428.80	40.52	7201.11	0.00	
c103C15	5	823.82	823.82	0.92	823.82	34.38	7200.18	0.00	823.82	73.45	7200.84	0.00	
c106C15	5	653.46	653.46	0.69	653.46	17.67	7200.19	0.00	653.46	63.86	7200.07	0.00	
c202C15	6	932.30	932.30	0.77	932.30	36.39	7200.23	0.00	932.30	68.58	7200.51	0.00	
c208C15	5	725.23	725.23	1.55	725.23	25.75	7200.17	0.00	725.23	68.69	7200.38	0.00	
r102C15	5	678.40	678.40	0.83	678.40	27.89	7200.17	0.00	678.40	64.94	7200.22	0.00	
r105C15	3	462.52	462.52	0.70	462.52	0.00	56.82	0.00	462.52	53.50	7200.10	0.00	
r202C15	3	528.59	535.08	1.41	528.59	30.25	7200.11	0.00	528.59	54.05	7200.95	0.00	
r209C15	2	369.29	371.60	1.26	369.29	7.10	7200.11	0.00	369.29	37.62	7201.49	0.00	
rc103C15	3	556.87	556.87	0.83	556.87	16.41	7200.07	0.00	556.87	58.12	7200.06	0.00	
rc108C15	3	510.41	511.03	1.19	510.41	3.47	7200.07	0.00	510.41	49.31	7200.14	0.00	
rc202C15	3	601.71	601.71	1.30	598.83	27.55	7200.18	-0.48	601.71	58.77	7200.24	0.00	
rc204C15	2	421.54	422.22	5.67	421.54	25.44	7201.01	0.00	421.54	49.29	7201.67	0.00	
Average		453.82	454.10	0.66	452.43	7.52	2803.97	-0.36	453.05	39.11	5707.60	-0.25	

using the MILP formulation with and without variable fixing and valid inequalities, i.e. constraints (2.27)-(2.39). For each instance, the MNS is run 10 times, and the solver is warm-started with the best solution from the 10 runs. A limit of 7200 s. is imposed on the solver. We observe that the MNS approach is stable across the instances and has a very small variance as expressed by the difference between the best and average result. The number of used vehicles ranges from one to six, with relatively more vehicles used for the clustered instances and those having shorter scheduling horizons, as expected.

Applied on the formulation with valid inequalities, Gurobi is able to solve with a proof of optimality all five-customer instances, all but one of the 10-customer instances and one 15-customer instance. For two other 15-customer instances, the MIP gap is brought to under 10%. For six instances, the solver is able to slightly improve the objective value given by the MNS. If we assume that optimality was reached within 7200 s., the MNS results are improved on average by 0.36%. Looking at the improvements, we can also observe that it is not necessarily its size alone that makes an instance challenging. Using the formulation without valid inequalities, Gurobi is only able to solve with a proof of optimality nine 5-customer instances within the solution time limit. The MIP gaps for the rest of the instances are significantly worse. Moreover, none of the MNS results of the 10-customer and 15-customer instances have been improved. This demonstrates that adding these problem-specific valid inequalities directly at the root node of the branching tree has a significant impact on the performance of the solver and thus serves as a better assessment of the corresponding performance of the MNS.

Table 2.3 also shows that the runtime of the solver, even when valid inequalities are used, grows exponentially with problem size. On the other hand, the runtime of the MNS is both shorter, even for the smallest instances, and grows in a much more stable manner. Since solution time is directly related to the neighborhood structure, in Table 2.4 we examine the added value of each neighborhood to the quality of the solution in terms of the mean best and mean average of the 36 instances over 10 runs. The first column reports the values obtained after the Initial Solution Construction (ISC) phase, while the rest of the columns show the improvement of these obtained by combining various neighborhoods.

Columns Nb-1, Nb-2 and Nb-3 demonstrate that the single most effective neighborhood for the modified Schneider et al. (2014) instances is reinsert, followed by 2-opt and swap. In all of

Table 2.4: Evaluation of MNS Neighborhoods and Their Combinations

	ISC Phase	Nb-1	Nb-2	Nb-3	Nb-12	Nb-13	Nb-23	Nb-123
Mean Best	529.12	-7.83%	-13.61%	-10.10%	-13.48%	-11.45%	-14.08%	-14.23%
Mean Avg.	529.12	-7.59%	-13.45%	-9.89%	-13.44%	-11.35%	-14.02%	-14.18%

Note. Nb-1: single- and inter-tour swap

Note. Nb-2: single- and inter-tour reinsert

Note. Nb-3: single- and inter-tour 2-opt

these neighborhood, as in the full MNS implementation, both single- and inter-tour moves are used. The next three columns demonstrate the gains of combining the neighborhoods. Not surprisingly, the two single best neighborhoods also produce the best combination. The last column represents the full implementation with all three neighborhoods and leads to a further visible improvement of the result compared to the best combination of two neighborhoods. Furthermore, during the experiments, we observed that the impact of including or not a neighborhood is not evenly spread across the instances, but affects certain ones more than others. Thus we are convinced that all three neighborhoods are beneficial for the performance of the MNS.

2.4.2 Tests on Benchmark Instances from the Literature

Table 2.5 presents our results on the MDVRPI (Crevier et al., 2007) instances. In all of them, the vehicles are stationed at a single depot, with the rest of the depots serving as intermediate facilities (IFs). There is a maximum tour duration, but no time windows or driver breaks. The instances are split into two sets. The first set (a1 to l1) contains 12 newly generated instances

Table 2.5: Comparison Against the BKS to the MDVRPI (Crevier et al., 2007) Instances

Inst- ance	(Cust., IFs)	Hemmelmayr et al. (2013)			MNS			Gap Best (%)	Gap Avg (%)
		Best	Avg	Runtime Avg (s.)	Best	Avg	Runtime Avg (s.)		
a1	(48,2)	1179.79	1180.57	85.20	1189.18	1202.89	21.12	0.80	1.89
b1	(96,2)	1217.07	1217.07	383.40	1217.07	1231.33	190.62	0.00	1.17
c1	(192,2)	1866.76	1867.96	1224.00	1885.57	1910.21	712.35	1.01	2.26
d1	(48,3)	1059.43	1059.43	94.20	1059.43	1071.19	19.33	0.00	1.11
e1	(96,3)	1309.12	1309.12	373.20	1309.12	1333.99	157.02	0.00	1.90
f1	(192,3)	1570.41	1573.05	1536.00	1576.81	1597.78	1148.62	0.41	1.57
g1	(72,4)	1181.13	1183.32	202.80	1186.59	1202.28	72.50	0.46	1.60
h1	(144,4)	1545.50	1548.61	876.60	1559.21	1571.26	531.82	0.89	1.46
i1	(216,4)	1922.18	1923.52	2014.80	1933.30	1956.97	1224.14	0.58	1.74
j1	(72,5)	1115.78	1115.78	166.80	1119.39	1139.20	66.34	0.32	2.10
k1	(144,5)	1576.36	1577.96	873.60	1581.23	1598.25	555.05	0.31	1.29
l1	(216,5)	1863.28	1869.70	2128.80	1880.93	1903.15	1435.59	0.95	1.79
a2	(48,4)	997.94	997.94	73.80	997.94	998.90	37.81	0.00	0.10
b2	(96,4)	1291.19	1291.19	384.60	1294.77	1343.87	217.86	0.28	4.08
c2	(144,4)	1715.60	1715.84	900.60	1731.60	1756.83	432.03	0.93	2.39
d2	(192,4)	1856.84	1860.92	1808.40	1863.97	1884.91	1031.17	0.38	1.29
e2	(240,4)	1919.38	1922.81	2958.60	1939.02	1979.30	1621.11	1.02	2.94
f2	(288,4)	2230.32	2233.43	4274.40	2273.17	2291.38	2451.33	1.92	2.59
g2	(72,6)	1152.92	1153.17	222.60	1153.21	1167.65	77.96	0.02	1.26
h2	(144,6)	1575.28	1575.28	939.60	1583.12	1601.21	506.46	0.50	1.65
i2	(216,6)	1919.74	1922.24	2515.20	1927.44	1958.01	1402.32	0.40	1.86
j2	(288,6)	2247.70	2250.21	4402.80	2259.99	2291.22	3056.50	0.55	1.82
Avg				1292.73			771.32	0.53	1.81

Chapter 2. The Waste Collection VRP

with two to five IFs and 48 to 216 customers. The second set (a2 to j2) contains 10 instances derived from those of Cordeau et al. (1997) by adding a central depot where the vehicles are stationed. The latter have four or six IFs and 48 to 288 customers. For both sets, the BKS are obtained by Hemmelmayr et al. (2013), who use a VNS with a dynamic programming procedure for the insertion of the intermediate facilities. Overall, in several cases we reach the BKS and our best solutions have an average gap of 0.53% with respect to the BKS. The gap with respect to the average solutions over 10 runs stands at 1.81%.

In order to assess the savings from allowing open tours with different origin and destination depots, we relax the MDVRPI instances by considering all intermediate facilities as possible destination depots of any vehicle. It should be noted that in the original MDVRPI formulation, intermediate facilities are actually depots with no vehicles stationed there. We consider the case of a relocation cost term weight ψ of zero in the objective function. Table 2.6 demonstrates that important savings can be obtained. Clearly, using a relocation cost term with a weight between zero and one would lead to results that fall between the restricted

Table 2.6: Savings from Allowing Open Tours in the MDVRPI (Crevier et al., 2007) Instances

Inst- ance	(Cust., IFs)	Hemmelmayr et al. (2013)			MNS			Gap Best (%)	Gap Avg (%)
		Best	Avg	Runtime Avg (s.)	Best	Avg	Runtime Avg (s.)		
a1	(48,2)	1179.79	1180.57	85.20	1119.46	1136.06	18.70	-5.11	-3.77
b1	(96,2)	1217.07	1217.07	383.40	1200.10	1206.99	176.90	-1.39	-0.83
c1	(192,2)	1866.76	1867.96	1224.00	1847.70	1884.80	857.52	-1.02	0.90
d1	(48,3)	1059.43	1059.43	94.20	1023.52	1040.32	22.95	-3.39	-1.80
e1	(96,3)	1309.12	1309.12	373.20	1293.14	1313.47	203.81	-1.22	0.33
f1	(192,3)	1570.41	1573.05	1536.00	1550.17	1589.25	1095.14	-1.29	1.03
g1	(72,4)	1181.13	1183.32	202.80	1142.31	1156.42	77.24	-3.29	-2.27
h1	(144,4)	1545.50	1548.61	876.60	1541.14	1559.50	518.71	-0.28	0.70
i1	(216,4)	1922.18	1923.52	2014.80	1895.89	1933.05	1420.72	-1.37	0.50
j1	(72,5)	1115.78	1115.78	166.80	1074.02	1082.40	70.05	-3.74	-2.99
k1	(144,5)	1576.36	1577.96	873.60	1553.86	1579.16	481.88	-1.43	0.08
l1	(216,5)	1863.28	1869.70	2128.80	1869.70	1894.59	1723.63	0.34	1.33
a2	(48,4)	997.94	997.94	73.80	911.82	923.37	38.01	-8.63	-7.47
b2	(96,4)	1291.19	1291.19	384.60	1263.30	1293.07	170.17	-2.16	0.15
c2	(144,4)	1715.60	1715.84	900.60	1694.27	1729.97	521.12	-1.24	0.82
d2	(192,4)	1856.84	1860.92	1808.40	1851.88	1870.52	984.40	-0.27	0.52
e2	(240,4)	1919.38	1922.81	2958.60	1927.62	1959.82	1794.59	0.43	1.92
f2	(288,4)	2230.32	2233.43	4274.40	2239.57	2282.11	2813.43	0.41	2.18
g2	(72,6)	1152.92	1153.17	222.60	1109.74	1132.96	76.15	-3.75	-1.75
h2	(144,6)	1575.28	1575.28	939.60	1573.16	1587.72	465.26	-0.13	0.79
i2	(216,6)	1919.74	1922.24	2515.20	1905.87	1924.26	1851.23	-0.72	0.11
j2	(288,6)	2247.70	2250.21	4402.80	2254.77	2284.15	3015.92	0.31	1.51
Avg				1292.73			836.25	-1.77	-0.37

Note. The gap from the BKS reflects the savings from allowing open tours when the home depots are as in the original instances.

2.4. Numerical Experiments

case and the one presented in Table 2.6. In the results we obtain, the savings over the restricted case are due to the fact that vehicles can now choose better destination depots, thus exploiting the geographical characteristics of the instances. It should be noted that all depots in these instances are centrally located. In a realistic situation where depots are not located in the center of the service area, but rather in a city's peripheral zones, the benefits are expected to be more pronounced.

Allowing the possibility of open tours is very important, but so is the choice of the actual home depot of each vehicle. According to the mathematical formulation presented in Section 2.2, starting from its home depot, a vehicle may be allowed to freely choose a destination depot. On the other hand, starting from any origin depot on a given day, the vehicle may be required to return to its home depot. Table 2.7 presents the case where we look for an optimal choice of the vehicles' home depots and also allow for open tours choosing the destination depots. We observe that savings grow significantly, in several cases reaching values in the order of

Table 2.7: Savings from Home Depot Optimization and Allowing Open Tours in the MDVRPI (Crevier et al., 2007) Instances

Inst- ance	(Cust., IFs)	Hemmelmayr et al. (2013)			MNS			Gap Best (%)	Gap Avg (%)
		Best	Avg	Runtime Avg (s.)	Best	Avg	Runtime Avg (s.)		
a1	(48,2)	1179.79	1180.57	85.20	1094.85	1106.46	20.52	-7.20	-6.28
b1	(96,2)	1217.07	1217.07	383.40	1208.23	1218.30	132.22	-0.73	0.10
c1	(192,2)	1866.76	1867.96	1224.00	1851.59	1885.82	764.14	-0.81	0.96
d1	(48,3)	1059.43	1059.43	94.20	1009.14	1023.26	27.05	-4.75	-3.41
e1	(96,3)	1309.12	1309.12	373.20	1280.14	1294.99	147.48	-2.21	-1.08
f1	(192,3)	1570.41	1573.05	1536.00	1544.27	1568.29	945.87	-1.66	-0.30
g1	(72,4)	1181.13	1183.32	202.80	1131.75	1138.56	65.15	-4.18	-3.78
h1	(144,4)	1545.50	1548.61	876.60	1523.97	1542.88	448.36	-1.39	-0.37
i1	(216,4)	1922.18	1923.52	2014.80	1900.70	1936.75	1443.26	-1.12	0.69
j1	(72,5)	1115.78	1115.78	166.80	1076.55	1080.02	68.83	-3.52	-3.21
k1	(144,5)	1576.36	1577.96	873.60	1525.45	1542.00	519.69	-3.23	-2.28
l1	(216,5)	1863.28	1869.70	2128.80	1846.76	1874.47	1249.17	-0.89	0.26
a2	(48,4)	997.94	997.94	73.80	887.58	911.09	45.10	-11.06	-8.70
b2	(96,4)	1291.19	1291.19	384.60	1256.27	1273.99	184.68	-2.70	-1.33
c2	(144,4)	1715.60	1715.84	900.60	1691.70	1715.05	421.25	-1.39	-0.05
d2	(192,4)	1856.84	1860.92	1808.40	1860.77	1870.70	833.90	0.21	0.53
e2	(240,4)	1919.38	1922.81	2958.60	1913.66	1951.95	2016.85	-0.30	1.52
f2	(288,4)	2230.32	2233.43	4274.40	2249.43	2274.70	2472.20	0.86	1.85
g2	(72,6)	1152.92	1153.17	222.60	1070.38	1085.91	119.28	-7.16	-5.83
h2	(144,6)	1575.28	1575.28	939.60	1550.94	1566.95	369.27	-1.54	-0.53
i2	(216,6)	1919.74	1922.24	2515.20	1903.29	1925.68	946.27	-0.86	0.18
j2	(288,6)	2247.70	2250.21	4402.80	2239.79	2271.01	2913.17	-0.35	0.92
Avg				1292.73			734.26	-2.54	-1.37

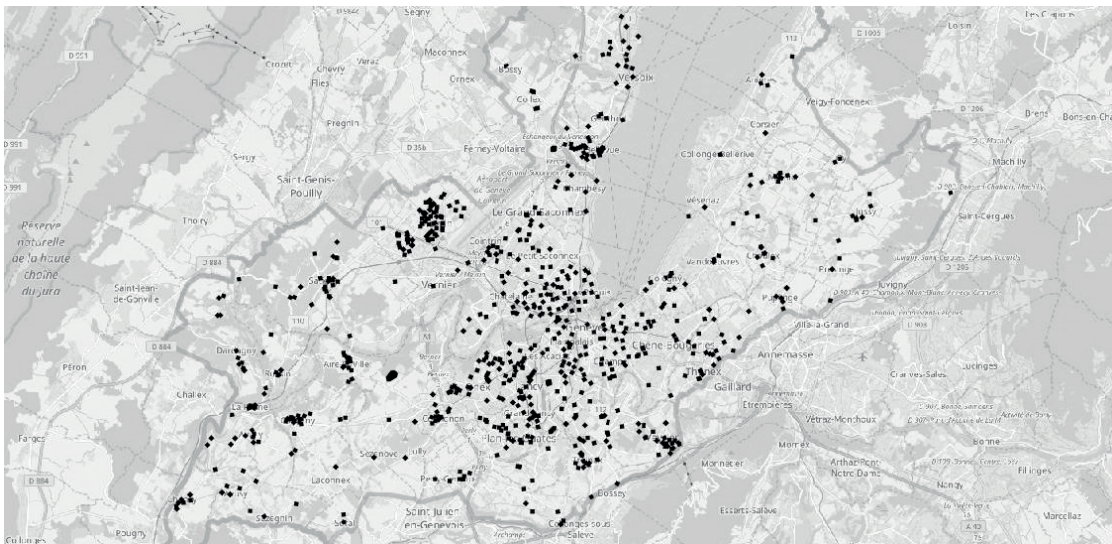
Note. The gap from the BKS reflects the savings from allowing open tours when the vehicles' home depots are also optimized.

10%. As expected, the savings for smaller instances with fewer depots and intermediate facilities are more pronounced, due to the fact that there are fewer intra-tour intermediate facility visits and thus the choice of destination depots is relatively more critical. Since real-world instances of our problem are of sizes comparable to the smaller Crevier et al. (2007) instances, these results are deemed an indication of the practical savings that can be obtained from this policy. The case study area to which this type of tours are applicable is in a French sparsely populated rural area. However, no historical tour data is available from the collector for comparison. The purpose of this test is therefore to justify the approach and quantify the benefits using synthetic instances.

2.4.3 Case Study

As mentioned in Section 2.4.2, there is no historical tour data available for the French region to which open tours are most applicable, which prompted us to evaluate their benefit on synthetic instances. In this section, we consider a case study of a collector of recyclable waste in the canton of Geneva, Switzerland. This collector uses specialized software for planning its collection tours and logging historical tour data. The software is currently used for planning the collection of white glass and PET¹. The available fleet consists of six vehicles with varying characteristics, whose weight capacities range from nine to 14 tons. The software can only plan one tour at a time and does not support all features required by the collector and present in our problem definition. Due to data confidentiality issues, we cannot disclose the complete information about the collection points and the collection process. Nevertheless, Figure 2.2 presents a map² of the collection points for recyclable

Figure 2.2: Geneva Service Area



¹Polyethylene terephthalate: a commonly used polymer for producing food and beverage containers.

²The map layer is from OpenStreetMap.


materials extracted from the cantonal open data portal (SITG, 2017). The area in question is 282.48 km² and has a population of approximately half a million. We remark that not all collection points are serviced by the collector that is used for this case study.

We obtained a sample of 35 planned tours for white glass and PET, their sizes ranging from seven to 38 containers, and with up to four dump visits per tour. Tour durations are rather short and the average vehicle speed is assumed to be 30 kmph. The distances between all depots, containers and dumps are shortest paths on a road network obtained from OpenStreetMap³. To perform a fair comparison between the software currently used for planning the tours and our MNS, we re-solve the problem for each tour separately, only enforcing the supported features, which are limited to the vehicles' volume and weight capacities. We keep the same origin and destination depot and provide all available dumps for the MNS to choose from. In the sample we obtained, all tours are planned for a different day, and could therefore use the same vehicle or visit the same container. As a consequence, we cannot combine multiple tours to be solved as a single instance. To compensate for the lack of richer features in the available real-world data, such features have been tested in the experiments in Sections 2.4.1 and 2.4.2.

Figure 2.3 compares the distances of the tours as planned by the software currently used by the collector and as provided by our MNS. The results of the MNS are averaged over 10 runs and computation times range from 0.05 to 7.58 s., with an average of 1.21 s. As the figure shows, all tours are improved and the average improvement per instance ranges from 1.73% to 34.91%, with a mean of 14.64%. It is interesting to observe that the distance improvements are due both to better container sequencing and better planned visits to the available recycling facilities.

After consultations with the concerned collector, we can estimate direct financial savings from fuel and labor in the order of 300,000 USD annually. These estimations assume that the number of tours is kept unchanged. However, given that the proposed solution approach

Figure 2.3: Distance Improvements Compared to the Currently Used Software



³<http://www.openstreetmap.org>.

can optimize multiple tours at the same time, rather than one at a time as in the current state of practice, further savings from a reduced number of tours, better planning of dump visits and more efficient labor utilization can also be expected. To give a better idea of the scale of the savings, we remark that the collection firm in question is one of several in an area of approximately half a million inhabitants.

It should be mentioned here that the software currently used by the collector requires a validation step once a planned tour has been executed. During validation, the collector deletes those containers from the originally proposed sequence that for some reason have not been collected when executing the tour. However, the ordering of the sequence cannot be changed. We use our MNS to build tours for the containers in the validated sequences and compare to their travel distances. The originally proposed sequences before validation are not available, but fortunately, not all validated tours have had containers removed, and those that have usually have one or two containers removed. Thus the results from Figure 2.3 provide strong evidence in favor of our MNS as compared to the current state of practice. Moreover, it solves a much richer problem and for a fraction of the computation time.

2.5 Summary

This chapter proposes a mathematical model with a number of valid inequalities for the waste collection VRP, an extension of the VRP-IF with a heterogeneous fixed fleet and the possibility of open tours with different origin and destination depots. The model includes several additional side constraints, such as time windows, a maximum tour duration, a mandated break period contingent on tour start time, multiple vehicle capacities, accessibility restrictions, and considers a general cost function corresponding to the cost structure of a typical firm. To solve realistic instances, we develop a meta-heuristic approach based on multiple neighborhood search.

The extensive computational testing confirms the advantage of including the valid inequalities in the optimization model. The MNS achieves optimality on small instances, exhibits competitive performance in comparison to state-of-the-art solution methods for special cases of our problem, and leads to important savings in the state of practice. We demonstrate that allowing open tours with different origin and destination depots can lead to noticeable savings especially in rural and sparsely populated areas where such benefits will be most pronounced. In addition, it presents fast computation times and outperforms significantly the solution currently in place in terms of quality and functionality.

The problem discussed here is already a very difficult to solve *NP*-hard problem. Chapter 3 integrates demand uncertainty into this problem, extending it to an inventory routing problem over a finite planning horizon. To solve it, we develop a more powerful and flexible algorithm—Adaptive Large Neighborhood Search (ALNS)—a state-of-the-art meta-heuristic approach with sophisticated search operators capable of tackling the added complexity.

3 The Waste Collection IRP with Stochastic Demands

This chapter is an extension of the technical report:

Markov, I., Bierlaire, M., Cordeau, J.-F., Maknoon, Y., and Varone, S. (2016). Inventory routing with non-stationary stochastic demands. Technical report TRANSP-OR 160825. Transport and Mobility Laboratory, EPFL, Lausanne, Switzerland.

The work therein has been performed by the author in collaboration with Prof. Michel Bierlaire, Prof. Jean-François Cordeau, Prof. Yousef Maknoon and Prof. Sacha Varone.

This chapter extends the waste collection VRP defined in Chapter 2 to an inventory routing problem over a planning horizon. In their survey of the IRP literature over the past thirty years, Coelho et al. (2014b) differentiate between finite and infinite planning horizons, with most problems modeled over finite horizons. Our waste collection IRP is also modeled as a finite-horizon problem. This is a natural choice given time discretization, which enters both at the demand forecasting and routing phases. Moreover, we solve a short-term operational problem which does not impose demand stationarity. Thus, a coarse approximation of the infinite future with the purpose of building a repetitive schedule would not be suitable for our problem.

Demand is the amount deposited in a container on a given day. It is stochastic, can be non-stationary, and is forecast using any model that provides the expected demands over the planning horizon and a measure of uncertainty represented by the standard deviation of the error terms, the latter assumed to be iid normal. Our waste collection IRP integrates demand uncertainty through the probabilities of container overflows and route failures. An overflow occurs when a container cannot accommodate further waste. In this case, the collector performs a recourse action by dispatching a vehicle on the same day to perform an emergency collection. This describes a back-ordering decision as demand is still served despite the overflow for the reason that people continue placing the waste beside the containers. An overflow also entails a fine by the municipality. All container collections follow the Order-

Up-to (OU) level inventory policy as containers are always fully emptied (Bertazzi et al., 2002). A route failure refers to an event where the vehicle runs out of capacity before the next scheduled dump visit due to higher than expected demand realizations of the collected containers (Dror and Trudeau, 1986). The recourse action is a visit to the nearest dump.

Generally speaking, our approach falls under the a priori optimization paradigm (Bertsimas et al., 1990; Gendreau et al., 2016), which takes a preventative attitude towards undesirable events and has the benefit of better tractability and solution consistency (Salavati-Khoshghalb et al., 2017). The system expects the undesirable events rather than being re-optimized after every occurrence. More precisely, the undesirable events are important in terms of their probabilities and approximate cost contributions. Given a rolling horizon approach, the expected cost of undesirable events and their recourse actions in the future periods of the planning horizon influences the here-and-now decisions, which are the ones we implement on a day-to-day basis. Finally, the problem discussed in this chapter has a single depot, with multiple depots reintroduced in Chapter 4 next. Given the multi-day planning horizon combined with uncertainty, we forgo the scheduling of driver breaks. The difficulty of this problem limits the use of fully exact approaches and has motivated the development of a state-of-the-art meta-heuristic algorithm based on Adaptive Large Neighborhood Search (ALNS).

The chapter is organized as follows. Section 3.1 positions our work with respect to the relevant literature. Section 3.2 outlines the forecasting model and formulates the stochastic IRP. Section 3.3 describes the ALNS algorithm. Section 3.4 presents the numerical experiments, and finally Section 3.5 ends with a summary of the main findings and contributions.

3.1 Related Literature

The vehicle routing subproblem embedded in our IRP already includes many rich routing features, notably a heterogeneous fixed fleet, time windows, a maximum tour duration, multiple dumps playing the role of intermediate facilities, accessibility restrictions, and a general cost function corresponding to the cost structure of a typical firm. The simultaneous presence of all these features is seldom considered in the VRP literature, and Chapter 2 already positions our contribution in this context. The problem we consider here has the complication of including these features in an IRP context. Thus, while they are essential to describing a realistic problem inspired from practice, they also pose a great challenge in terms of modeling and solution methodology. The IRP is an *NP*-hard optimization problem which decides simultaneously the vehicle tours, the visit days, and as a consequence the collection quantities. Comprehensive surveys on the IRP can be found in Abdelmaguid (2004), Moin and Salhi (2007), Andersson et al. (2010), Yu et al. (2012), Bertazzi and Speranza (2013), Coelho et al. (2014b), Ivarsøy and Solhaug (2014) and Park et al. (2016), and a particular focus on stochastic problems and aspects can be found in Moin and Salhi (2007), Yu et al. (2012), Bertazzi and Speranza (2013) and Coelho et al. (2014b). In the following, we limit

our attention to finite-horizon stochastic problems, i.e. the class to which our problem belongs. We also put particular focus on the importance of the rolling horizon approach (see Section 1.1 in Chapter 1).

Trudeau and Dror (1992) extend the work of Dror and Ball (1987) on the optimal service frequency under a stochastic setting. They consider both stock-outs and route failures. Unlike previous research (see e.g. Stewart and Golden, 1983; Dror et al., 1985; Dror and Levy, 1986; Dror and Ball, 1987; Larson, 1988) which uses a vehicle with an artificially small capacity to avoid route failures, Trudeau and Dror (1992) develop an analytical probability expression, and corroborate their modeling approach with a simulation experiment. Our work differs from that of Trudeau and Dror (1992) in several major aspects. In particular, we have a heterogeneous fixed fleet. Route failures in our case apply to depot-to-dump or dump-to-dump trips, of which there could be several in a given tour. Finally, we do not impose a maximum of one visit and one overflow per container during the planning horizon, which precludes the derivation of exact closed-form probability measures. On the contrary, it requires the complicated management of binary trees, tracking each container's visit-dependent and conditional probabilities of overflow on each day of the planning horizon. In addition, we consider multiple rich routing features.

The work of Bard et al. (1998b) includes intermediate facilities in a distribution context. They apply problem decomposition with a two-week rolling horizon. Customers to be visited during the planning horizon are identified and those scheduled for the first week are routed, after which the horizon is rolled over by a week. The customer selection procedure is based on Jaillet et al. (2002) who derive the optimal restocking frequency and the incremental cost of deviating from it. In the first step of the decomposition scheme, customers whose optimal visit day falls within the two-week horizon are assigned to specific days by solving a balanced generalized assignment problem that minimizes the total incremental cost, accounting for uncertainty through a lower and upper bound on the total daily demand to be served. The solution of the routing problem relies on construction and improvement heuristics including inter-day customer exchanges. Similar ideas, based on the identification of customers who must be served versus those who may be served are used in Bitsch (2012) and Mes et al. (2014), both with applications to waste collection where the objective is the minimization of overflows. The former relies on the calculation of incremental costs, while the latter on expectation-based service frequency. Due to the implied repetitive pattern, this type of approaches is only appropriate in situations where demand stationarity can be assumed. Our problem does not impose this restriction.

Campbell and Savelsbergh (2004) also deal with uncertainty through a decomposition approach that solves the problem of assigning customers to days first, using the cost of a giant TSP tour as a crude measure of the daily routing cost, and with coarser period aggregations toward the end of the planning horizon. Afterwards, the IRP is solved for the first few days of the planning horizon for the customers that were assigned there and assuming deterministic information. This approach is used in a rolling horizon framework with the

benefit of reflecting longer-term costs in the shorter-term problem, i.e. on the days for which the actual IRP is solved. Such a balance, usually expressed through a so-called reduction procedure, was the focus of much of the above-mentioned IRP research (see Dror and Ball, 1987; Trudeau and Dror, 1992; Dror and Trudeau, 1996; Jaillet et al., 2002). Stochasticity is also discussed in Coelho et al. (2014a), who present a modeling and solution framework for dynamic and stochastic IRP, incorporating the use of forecasting. However, their approach relies on constructing point forecasts to be used in a rolling horizon fashion without the explicit treatment of uncertainty. Independent of the modeling approach or the methodology used, the rolling horizon technique is useful in dealing with uncertainty by helping make forward-looking decisions in the operational short-term.

More recently, research on the IRP has dealt with uncertainty in various ways. Solyalı et al. (2012), for example, use the robust optimization approach introduced by Bertsimas and Sim (2003, 2004) to solve a problem with dynamic uncertain demands, ensuring that vehicle capacity is not violated for any realization of the customer demands. They develop a strong formulation and use a branch-and-cut solution approach. Bertazzi et al. (2013) propose a heuristic rollout algorithm that uses a sampling approach to generate demand scenarios for the current period and considers the average demand for future ones. Decisions are made by solving a mixed integer program by branch-and-cut in each period. A similar approach is used by Bertazzi et al. (2015) who apply it to an IRP with transportation procurement. Adulyasak et al. (2015) propose a two-stage and a multi-stage approach for a production-routing problem under demand uncertainty, in which the first stage determines production setup and visit frequencies, while subsequent stages determine production and delivery quantities. They develop exact formulations and a branch-and-cut algorithm, and a Benders decomposition approach able to solve instances of realistic size for a high number of scenarios. Stochastic optimization with recourse is used by Hemmelmayr et al. (2010) and Nolz et al. (2014b), who present applications related to blood product distribution and medical waste collection, respectively. Chance-constrained approaches, often oriented towards maintaining a service level, can be found in Yu et al. (2012), Abdollahi et al. (2014), Soysal et al. (2015) and Soysal et al. (2016), while static risk expressions in the objective function that use the demand distribution parameters are applied by Nekooghadirli et al. (2014a) and Nekooghadirli et al. (2014b). Ribeiro and Lourenço (2003) apply recursive logic to integrate inventory cost in the objective function for exponentially distributed demands and propose a simple heuristic approach to tackle their problem.

Chapter 1 discussed the above approaches along with their advantages and limitations, and identified our goal of pricing demand uncertainty as would a typical cost minimizing firm. Our modeling approach uses stochastic information in the objective function, integrating the probabilities of container overflows and route failures. We consider both the cost of these undesirable events and their associated recourse actions. We can pre-compute the probabilities of container overflows, while those of route failure can be approximated precisely at runtime. This leads to a tractable approach capable of treating a variety of rich routing features seldom considered in the IRP literature.

3.2 Formulation

In what follows, Section 3.2.1 presents a brief sketch of the forecasting model, Section 3.2.2 develops a mathematical formulation for our stochastic IRP (SIRP), and Section 3.2.3 discusses the necessary changes to the latter for solving benchmark instances from the literature. Table 3.1 summarizes the used notations. Some of the notations, in particular the inventory holding cost, are only used in the model reformulations presented in Section 3.2.3 but are still included in the table for completeness and ease of reference. We note that container demand refers to the volume amount placed in a container on a given day. Container inventory and capacity are also measured in terms of volume. Vehicles, on the other hand, have both volume and weight capacities. Depending on the density of the collected waste flow,

Table 3.1: Notations

Sets			
\mathcal{V}	set of distinct container deposit volumes	\mathcal{H}	historical estimation period
\mathcal{T}	planning horizon = $\{0, \dots, u\}$	\mathcal{T}^+	shifted planning horizon = $\{1, \dots, u, u + 1\}$
o	origin	d	destination
\mathcal{P}	set of containers	\mathcal{D}	set of dumps
\mathcal{N}	set of all points = $\{o\} \cup \{d\} \cup \mathcal{P} \cup \mathcal{D}$	\mathcal{K}	set of vehicles
\mathcal{S}_{kt}	set of depot-to-dump or dump-to-dump trips for vehicle $k \in \mathcal{K}$ on day $t \in \mathcal{T}$	\mathcal{S}	set of containers in a particular trip in \mathcal{S}_{kt}
Parameters			
l_{itg}	Poisson rate for deposit volume v of container i on day t		
κ_{it}	vector of covariates for container i on day t		
ϖ_v	vector of estimable parameters for deposit volume v		
φ_k	daily deployment cost of vehicle k (monetary)		
β_k	unit-distance running cost of vehicle k (monetary)		
θ_k	unit-time running cost of vehicle k (monetary)		
Ω_k	capacity of vehicle k		
π_{ij}	travel distance of arc (i, j)		
τ_{ijk}	travel time of vehicle k on arc (i, j)		
λ_i, μ_i	lower and upper time window bound at point i		
δ_i	service duration at point i		
ω_i	capacity of container i		
η_i	inventory holding cost at point i (monetary)		
a_{kt}	1 if vehicle k is available on day t , 0 otherwise		
a_{ik}	1 if container i is accessible by vehicle k , 0 otherwise		
ρ_{it}	demand of container i on day t		
ε_{it}	error term of container i on day t		
ζ	forecasting error (standard deviation of the error terms)		
σ_{it}	1 indicates that container i is in a state of full and overflowing on day t , 0 otherwise		
χ	container overflow cost (monetary)		
ζ	container emergency collection cost (monetary)		
H	maximum tour duration		
ψ	Route Failure Cost Multiplier (RFCM) $\in [0, 1]$		
C_S	the average routing cost of going from $\mathcal{S} \in \mathcal{S}_{kt}$ to the nearest dump and back to \mathcal{S} (monetary)		
Decision Variables			
$x_{ijk t}$	1 if vehicle k traverses arc (i, j) on day t , 0 otherwise (binary)		
$y_{ik t}$	1 if vehicle k visits point i on day t , 0 otherwise (binary)		
$z_{k t}$	1 if vehicle k is used on day t , 0 otherwise (binary)		
$q_{ik t}$	expected pickup quantity by vehicle k from container i on day t (continuous)		
$Q_{ik t}$	expected cumulative quantity on vehicle k at point i on day t (continuous)		
I_{it}	expected inventory of container i at the start of day t (continuous)		
$S_{ik t}$	start-of-service time of vehicle k at point i on day t (continuous)		

one of them becomes limiting while the other may not be. However, we observe that if the weight capacity becomes limiting before the volume capacity, the volume capacity can be adjusted to become limiting at the same time. Through this simple preprocessing step, we avoid tracking both volume and weight, thus simplifying the notation of Chapter 2.

3.2.1 Forecasting Model

To forecast the expected container demands over the planning horizon and to derive the forecasting error, we use the forecasting model proposed by Markov et al. (2015), which exhibits superior in- and out-of-sample performance compared to alternatives. It is based on a discrete mixture of count-data models describing populations depositing different waste volumes in the containers. Thus, it supposedly captures a realistic underlying behavior though simplified. We assume a set \mathcal{V} of distinct deposit volumes, where deposit volume $v \in \mathcal{V}$ is generated with a Poisson rate l_{itv} for container i on day t . The rate l_{itv} takes the functional form $l_{itv} = \exp(\boldsymbol{\kappa}_{it}^\top \boldsymbol{\omega}_v)$, where $\boldsymbol{\kappa}_{it}$ is a vector of covariates, such as the day of the week, weather variables, holiday periods, etc., and $\boldsymbol{\omega}_v$ is a vector of estimable parameters for deposit volume v . We formulate an expression for the expected value of the demand of container i on day t as follows:

$$\mathbb{E}(\rho_{it}) = \sum_{v \in \mathcal{V}} v l_{itv}. \quad (3.1)$$

To fit the model, we minimize the sum of squared errors between the observed ρ_{it}^o and the expected demand $\mathbb{E}(\rho_{it})$ over the set of containers \mathcal{P} and a historical period \mathcal{H} of data availability:

$$\min \sum_{i \in \mathcal{P}} \sum_{t \in \mathcal{H}} \left(\rho_{it}^o - \sum_{v \in \mathcal{V}} v l_{itv} \right)^2, \quad (3.2)$$

assuming strict exogeneity and with error terms represented by white noise as:

$$\rho_{it} = \mathbb{E}(\rho_{it}) + \varepsilon_{it}, \quad \text{where } \varepsilon_{it} \text{ are iid normal,} \quad (3.3)$$

and where a consistent estimate of the variance is given by:

$$\zeta^2 = \frac{\sum_{i \in \mathcal{P}} \sum_{t \in \mathcal{H}} (\rho_{it}^o - \mathbb{E}(\rho_{it}))^2}{|\mathcal{P}||\mathcal{H}| - \#\text{params}}. \quad (3.4)$$

We refer to ζ as the forecasting error. The denominator in formula (3.4) is the total number of data observations $|\mathcal{P}||\mathcal{H}|$ minus the number of estimated parameters in the model. For a more detailed description of the model, the reader is referred to Markov et al. (2015). Chapter 4 discusses the relaxation of the iid normality assumption of the error terms in equation (3.3).

3.2.2 Stochastic IRP Model

Our SIRP is defined for a planning horizon $\mathcal{T} = \{0, \dots, u\}$ and we are given a complete directed graph $\mathcal{G}(\mathcal{N}, \mathcal{A})$, with $\mathcal{N} = \{o\} \cup \{d\} \cup \mathcal{P} \cup \mathcal{D}$, where o and d represent the depot as an origin and a destination, respectively, \mathcal{P} is the set of containers, \mathcal{D} is the set of dumps, and $\mathcal{A} = \{(i, j) : \forall i, j \in \mathcal{N}, i \neq j\}$ is the set of arcs. For modeling purposes, it is assumed that the set \mathcal{D} contains a sufficient number of replications of each dump to allow multiple visits by the same vehicle on the same day. Multiple depots are reintroduced in Chapter 4.

There is an asymmetric distance matrix, with π_{ij} the travel distance of arc (i, j) . Each vehicle may have a different average speed, which results in a vehicle-specific travel time matrix, where τ_{ijk} is the travel time of vehicle k on arc (i, j) . Each point has a single time window $[\lambda_i, \mu_i]$, where λ_i and μ_i stand for the earliest and latest possible start-of-service time. Start of service after μ_i is not allowed, and if the vehicle arrives before λ_i , it has to wait. Service duration at each point is denoted by δ_i . For containers it is mostly influenced by the type of container, e.g. underground or overground, and for dumps by factors such as weighing and billing. Hence service duration is not indexed by vehicle. Service duration at the depots is zero. There is an expected demand $\mathbb{E}(\rho_{it})$ for container i on day t . Container capacity is denoted by ω_i , and a cost χ is charged for a full and overflowing container. There is a heterogeneous fixed fleet \mathcal{K} , with each vehicle defined by its capacity Ω_k , a daily deployment cost φ_k , a unit-distance running cost β_k , and a unit-time running cost θ_k . The binary flags a_{kt} denote whether vehicle k is available on day t , and the binary flags α_{ik} denote whether container i is accessible by vehicle k . The maximum tour duration is denoted by H .

We introduce the following binary decision variables: $x_{ijk_t} = 1$ if vehicle k traverses arc (i, j) on day t , 0 otherwise; $y_{ikt} = 1$ if vehicle k visits point i on day t , 0 otherwise; $z_{kt} = 1$ if vehicle k is used on day t , 0 otherwise. In addition, the following continuous variables are used: q_{ikt} for the expected pickup quantity by vehicle k from container i on day t ; Q_{ikt} for the expected cumulative quantity on vehicle k arriving at point i on day t ; I_{it} for the expected inventory of container i at the start of day t ; and S_{ikt} for the start-of-service time of vehicle k at point i on day t . The inventory levels at the start of the planning horizon I_{i0} are known with certainty. For modeling purposes, we assume that container inventory is updated at the start of each day before vehicle visits. This simplification is necessary due to the daily time discretization. As a consequence, the pickup quantity is independent of the time of day that the vehicle collects a container.

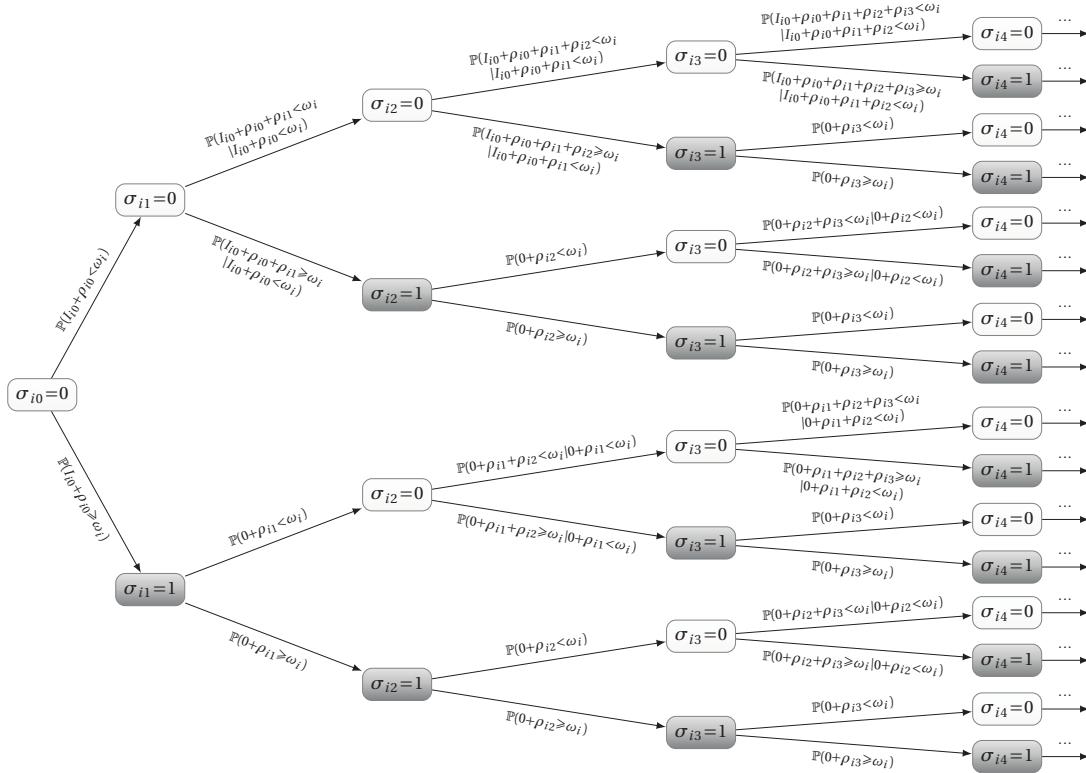
Derivation of the Overflow Probabilities.

Unlike in most traditional IRPs, we have no inventory holding costs at the containers or dumps. To formulate the objective function, we introduce the notions of a regular and an emergency collection. Let σ_{it} denote the state of container i on day t , where $\sigma_{it} = 0$ denotes that container i is not full on day t , while $\sigma_{it} = 1$ denotes that it is full and overflowing. A regular collection of container i on day t by vehicle k is one for which $y_{ikt} = 1$. On

the other hand, an emergency collection occurs when the container is in a state $\sigma_{it} = 1$ and $y_{ikt} = 0, \forall k \in \mathcal{K}$. An emergency collection incurs a high cost ζ , which is an approach often employed in the IRP literature (e.g. Dror and Ball, 1987; Trudeau and Dror, 1992; Hemmelmayr et al., 2010; Coelho et al., 2014a), and empties the container in question. Our routing cost is thus counterbalanced by the container overflow cost χ and the emergency collection cost ζ .

Let us start with an initial inventory I_{i0} such that container i is initially in state $\sigma_{i0} = 0$. If the container never undergoes a regular collection during the planning horizon, its state probability tree develops as illustrated in Figure 3.1. We observe that all branches starting from a state $\sigma_{it} = 0$ involve the calculation of conditional probabilities, while those starting from a state $\sigma_{it} = 1$ involve unconditional probabilities because the inventory is set to zero by the emergency collection. For our problem, we are only interested in the probability of overflow, i.e. of being in a state $\sigma_{it} = 1$. For day $t = 0$, this is either 0 or 1, depending on the initial state, while for all other days it is obtained by successively multiplying the branch probabilities. If we impose a regular collection on day $t = 2$, the probability of overflow on day $t = 2$ is the probability of being in state $\sigma_{i2} = 1$. To calculate the probability of overflow for subsequent days, we start a new tree with a root on day $t = 2$. Without loss of generality, we can set the root of the new tree to state $\sigma_{i2} = 1$ since the inventory is set to zero by

Figure 3.1: Container State Probability Tree



the regular collection and the two initial branches of the new tree will have unconditional probabilities. Regardless of the initial state, all branch probabilities can be precomputed, including those that occur when a new tree is started by a regular collection. For container i , the exhaustive list is given by:

- The unconditional probability of overflow with non-zero initial inventory. This only applies at the root node of the state probability tree on day $t = 0$ and is given by $\mathbb{P}(I_{i0} + \rho_{i0} \geq \omega_i)$.
- The unconditional probabilities of overflow with zero initial inventory. These apply either at the root node or at a state of overflow and are expressed by $\mathbb{P}(0 + \rho_{ih} \geq \omega_i), \forall h \in \mathcal{T}$.
- The conditional probabilities of overflow with non-zero initial inventory. These apply along the tree's uppermost branch and write as $\mathbb{P}(I_{i0} + \sum_{t=0}^h \rho_{it} \geq \omega_i \mid I_{i0} + \sum_{t=0}^{h-1} \rho_{it} < \omega_i), \forall h \in \mathcal{T} : h > 0$.
- The conditional probabilities of overflow with zero initial inventory apply in all other cases and are obtained as $\mathbb{P}(0 + \sum_{t=g}^h \rho_{it} \geq \omega_i \mid 0 + \sum_{t=g}^{h-1} \rho_{it} < \omega_i), \forall g, h \in \mathcal{T} : h > g$.

The calculation of the conditional probabilities involves the evaluation of:

$$\mathbb{P}\left(I_{ig} + \sum_{t=g}^h \rho_{it} \geq \omega_i \mid I_{ig} + \sum_{t=g}^{h-1} \rho_{it} < \omega_i\right). \quad (3.5)$$

Given the definition of the error terms as iid normal in equation (3.3), expression (3.5) takes the form:

$$\mathbb{P}\left(\sum_{t=g}^h \varepsilon_{it} \geq \omega_i - I_{ig} - \sum_{t=g}^h \mathbb{E}(\rho_{it}) \mid \sum_{t=g}^{h-1} \varepsilon_{it} < \omega_i - I_{ig} - \sum_{t=g}^{h-1} \mathbb{E}(\rho_{it})\right). \quad (3.6)$$

Substitute $a = \omega_i - I_{ig} - \sum_{t=g}^{h-1} \mathbb{E}(\rho_{it})$, and $X = \sum_{t=g}^{h-1} \varepsilon_{it}$, where $X \sim \mathcal{N}(0, (h-g)\zeta^2)$ and X is independent of ε_{ih} . Formula (3.6) then rewrites as:

$$\begin{aligned} \mathbb{P}(X + \varepsilon_{ih} \geq a - \mathbb{E}(\rho_{ih}) \mid X < a) &= \frac{\mathbb{P}(\varepsilon_{ih} \geq a - \mathbb{E}(\rho_{ih}) - X, X < a)}{\mathbb{P}(X < a)} = \\ &= \frac{1}{\Phi_X(a)} \times \frac{1}{2\pi\zeta^2\sqrt{h-g}} \int_{-\infty}^a \int_{a-\mathbb{E}(\rho_{ih})-x}^{\infty} e^{-\frac{x^2}{2(h-g)\zeta^2}} e^{-\frac{y^2}{2\zeta^2}} dx dy, \end{aligned} \quad (3.7)$$

where $\Phi_X(\cdot)$ is the CDF of X . We standardize the joint probability in expression (3.7) by setting $x = x/(\zeta\sqrt{h-g})$ and $y = y/\zeta$, and thus arrive at expression (3.8) for the conditional

probability we are looking for:

$$\begin{aligned} \mathbb{P}(X + \varepsilon_{ih} \geq a - \mathbb{E}(\rho_{ih}) \mid X < a) &= \frac{1}{2\pi\Phi\left(\frac{a}{\zeta\sqrt{h-g}}\right)} \int_{-\infty}^{\frac{a}{\zeta\sqrt{h-g}}} \int_{\frac{a-\mathbb{E}(\rho_{ih})-x\zeta\sqrt{h-g}}{\zeta}}^{\infty} e^{-\frac{x^2}{2}} e^{-\frac{y^2}{2}} dx dy = \\ &= \frac{1}{2\sqrt{2}\pi\Phi\left(\frac{a}{\zeta\sqrt{h-g}}\right)} \int_{-\infty}^{\frac{a}{\zeta\sqrt{h-g}}} e^{-\frac{x^2}{2}} \operatorname{erfc}\left(\frac{a-\mathbb{E}(\rho_{ih})-x\zeta\sqrt{h-g}}{\zeta\sqrt{2}}\right) dx, \end{aligned} \quad (3.8)$$

where $\Phi(\cdot)$ is the CDF of a standard normal random variable. The single integral in expression (3.8) can be evaluated using a standard statistical package like R in the order of milliseconds. For a problem of realistic size, all the necessary unconditional and conditional probabilities can be automatically precomputed in a negligible amount of time using the latest container information. The derivations above, presented here in the context of the waste collection IRP, are extended to a general finite horizon routing context under mild assumptions in Chapter 4.

Objective Function.

We are now in a position to formulate the objective function z which comprises the Routing Cost (RC), the Expected Overflow and Emergency Collection Cost (EOECC), and the Expected Route Failure Cost (ERFC):

$$\min z = \text{RC} + \text{EOECC} + \text{ERFC}. \quad (3.9)$$

The routing cost reflects the daily deployment cost, the distance-related cost and the time-related cost for each used vehicle over the planning horizon. It is formulated as:

$$\text{RC} = \sum_{t \in \mathcal{T}} \sum_{k \in \mathcal{K}} \left(\varphi_k z_{kt} + \beta_k \sum_{i \in \mathcal{N}} \sum_{j \in \mathcal{N}} \pi_{ij} x_{ijk t} + \theta_k (S_{dkt} - S_{okt}) \right). \quad (3.10)$$

The expected overflow and emergency collection cost is expressed as:

$$\begin{aligned} \text{EOECC} = \sum_{t \in \mathcal{T} \cup \mathcal{T}^+} \sum_{i \in \mathcal{P}} &\left(\mathbb{P}(\sigma_{it} = 1 \mid m = \max(0, g \in \mathcal{T} : g < t : \exists k \in \mathcal{K} : y_{ikg} = 1)) \right. \\ &\left. \times \left(\chi + \zeta - \zeta \sum_{k \in \mathcal{K}} y_{ikt} \right) \right), \end{aligned} \quad (3.11)$$

where the probability of being in a state of overflow is conditional on the most recent regular collection, identified for each container i by the index m . For a given container i , the max operator returns the day g of the most recent regular collection, or 0 if the container has not undergone any regular collections before day t . The state probability is calculated by multiplication of the involved branch probabilities on the tree, where conditional probabilities

are computed using formula (3.5) and the discussed methodology (3.6)–(3.8). For a day t , the applied cost includes the container overflow cost χ and the emergency collection cost ζ in case there is no regular collection on that day, and only the container overflow cost χ in case there is a regular collection. Although there is no uncertainty on day $t = 0$, we still need to pay the overflow cost if the container is in a state of overflow. On the other hand, the inventories at the start of the first day after the end of the planning horizon are completely determined by the decisions taken during the planning horizon. For this reason, the EOEC is computed for $t \in \mathcal{T} \cup \mathcal{T}^+$, where $\mathcal{T}^+ = \{1, \dots, u, u + 1\}$ is the planning horizon shifted right by one day.

The expected route failure cost reflects the vehicles' inability due to insufficient capacity to serve the containers on the scheduled depot-to-dump or dump-to-dump trips. It is expressed as:

$$\text{ERFC} = \sum_{t \in \mathcal{T} \setminus \{0\}} \sum_{k \in \mathcal{K}} \sum_{\mathcal{S} \in \mathcal{S}_{kt}} \left(\psi C_{\mathcal{S}} \mathbb{P} \left(\sum_{s \in \mathcal{S}} \left(\Lambda_{sm} + \sum_{h=m}^{t-1} \rho_{sh} \right) > \Omega_k \mid m = \max(0, g \in \mathcal{T} : g < t : \exists k' \in \mathcal{K} : y_{sk'g} = 1) \right) \right) \quad (3.12)$$

where \mathcal{S}_{kt} is the set of depot-to-dump or dump-to-dump trips for vehicle k on day t , \mathcal{S} is the set of containers in a particular trip, $C_{\mathcal{S}}$ is the average routing cost of going from this set to the nearest dump and back, and Λ_{sm} is the inventory of container s after regular collection on day m . The set \mathcal{S}_{kt} is generated by inspecting the routing variables x_{ijk_t} . At every feasible solution, for each vehicle k on each day t we can inspect the point visit sequence encoded in the variables x_{ijk_t} to generate the set of depot-to-dump and dump-to-dump trips. The parameter $\psi \in [0, 1]$, which we refer to as the Route Failure Cost Multiplier (RFCM), is used to scale up or down the degree of conservatism regarding this cost component. The probability is conditional on the most recent regular collection identified for each container s by the index m . For a given container s , the max operator returns the day g of the most recent regular collection, or 0 if the container has not undergone any regular collections before day t . The inventory of container s after regular collection on day m is defined as:

$$\Lambda_{sm} = I_{sm} - \sum_{k \in \mathcal{K}} q_{skm}. \quad (3.13)$$

Given an order-up-to policy, this value is 0 if there is a regular collection on day m , and is equal to the initial inventory I_{s0} if there is no regular collection on day 0. In essence, the probability of a route failure in a set \mathcal{S} is the probability that the sum of the random daily demands, plus potentially the initial inventories on day 0, collected from this set exceeds the vehicle capacity. By definition, there are no route failures on day $t = 0$ as the container information is fully known.

The nearest dump to each container can be precomputed. Probability-wise, once the days

m of the previous collection of each container are found, the remaining probability is unconditional. It is impractical to precompute for all possible combinations on the left-hand side of the probability formula in the ERFC expression (3.12). Thus, we implement a solution in which the probability is evaluated during runtime using an approximation of the standard normal distribution based on the approximation of the error function:

$$\operatorname{erf}(x) \approx 1 - (a_1 t + a_2 t^2 + \dots + a_5 t^5) e^{-x^2}, \quad t = \frac{1}{1 + dx}, \quad (3.14)$$

whose parameter values are $d = 0.3275911$, $a_1 = 0.254829592$, $a_2 = -0.284496736$, $a_3 = 1.421413741$, $a_4 = -1.453152027$, and $a_5 = 1.061405429$, and whose maximum approximation error is 1.5×10^{-7} (Abramowitz and Stegun, 1972). These repetitive calculations have no discernible influence on the algorithm's runtime. Chapter 4 discusses the complications in calculating the route failure probabilities that arise from relaxing the iid normal assumption (3.3).

The objective function is non-linear due to the non-linear nature of the EOECC and the ERFC components defined above. On the other hand, the formulations of the RC and the ERFC ignore the probability of containers overflowing before the day t on which they are collected. Skipping such containers in the tours performed on day t would reduce the RC. It would also reduce the ERFC due to the lower probability of the collected volume on day t exceeding the vehicle capacity. Modeling these two effects without imposing additional assumptions would make the probability expressions intractable. On the other hand, given the operational nature of the problem, developing a full-blown simulator of the objective function to capture them would be very impractical. As a result, our objective function is an overestimation of the real cost. We revisit this question in Chapter 4, where we perform simulation-validation experiments that demonstrate that the level of overestimation is insignificant, and thus our objective function is an excellent representation of the real cost.

Constraints.

The constraints are extended and adapted from those presented Section 2.2.1 in Chapter 2 and can be split into several categories, with the first category consisting of basic vehicle routing constraints. Constraints (3.15) and (3.16) ensure that only available vehicles are used, and that if a vehicle is used, its tour starts at the origin and ends at the destination, with a visit to a dump immediately before that. Constraints (3.17) link the visit and the routing variables, while constraints (3.18) stipulate that a container is visited by at most one vehicle on a given day. Constraints (3.19) guarantee that vehicles do not visit inaccessible points. Flow conservation is represented by constraints (3.20).

$$\sum_{j \in \mathcal{N}} x_{ojkt} = \alpha_{kt} z_{kt}, \quad \forall t \in \mathcal{T}, k \in \mathcal{K} \quad (3.15)$$

$$\sum_{i \in \mathcal{D}} x_{idkt} = \alpha_{kt} z_{kt}, \quad \forall t \in \mathcal{T}, k \in \mathcal{K} \quad (3.16)$$

$$y_{ikt} = \sum_{j \in \mathcal{N}} x_{ijkt}, \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, i \in \mathcal{P} \quad (3.17)$$

$$\sum_{k \in \mathcal{K}} y_{ikt} \leq 1, \quad \forall t \in \mathcal{T}, i \in \mathcal{P} \quad (3.18)$$

$$y_{ikt} \leq \alpha_{ik}, \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, i \in \mathcal{P} \quad (3.19)$$

$$\sum_{i \in \mathcal{N}} x_{ijk} = \sum_{i \in \mathcal{N}} x_{jik}, \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, j \in \mathcal{P} \cup \mathcal{D} \quad (3.20)$$

The inventory constraints are necessary for tracking the container inventories and linking them to the vehicle visits and the pickup quantities. Constraints (3.21) track the inventories as a function of the previous day's inventories, pickup quantities and expected demands. Constraints (3.22) impose the fact that, in expected terms, we do not accept container overflows. As already mentioned in Section 3.2.2, the inventories need to be computed over \mathcal{T}^+ , starting from the fully known inventories on day $t = 0$. Constraints (3.23) ensure that if the starting inventory exceeds capacity, the container must be collected on day $t = 0$. The big- M reflects the assumption that the expected daily demand can never exceed the container capacity. In addition, a daily rolling horizon enforces the one-day back-order limit. Constraints (3.24) force the pickup quantity to zero if the container is not visited. Constraints (3.25) and (3.26) represent the order-up-to policy. The big- M values in constraints (3.24) and (3.26) can be set to $2\omega_i$ for $t = 0$ and to ω_i otherwise, reflecting the fact that the expected pick-up quantity cannot exceed container capacity, except on day $t = 0$.

$$I_{it} = I_{i(t-1)} - \sum_{k \in \mathcal{K}} q_{ik(t-1)} + \mathbb{E}(\rho_{i(t-1)}), \quad \forall t \in \mathcal{T}^+, i \in \mathcal{P} \quad (3.21)$$

$$I_{it} \leq \omega_i, \quad \forall t \in \mathcal{T}^+, i \in \mathcal{P} \quad (3.22)$$

$$I_{i0} - \omega_i \leq \omega_i \sum_{k \in \mathcal{K}} y_{ik0}, \quad \forall i \in \mathcal{P} \quad (3.23)$$

$$q_{ikt} \leq M y_{ikt}, \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, i \in \mathcal{P} \quad (3.24)$$

$$q_{ikt} \leq I_{it}, \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, i \in \mathcal{P} \quad (3.25)$$

$$q_{ikt} \geq I_{it} - M(1 - y_{ikt}), \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, i \in \mathcal{P} \quad (3.26)$$

In the context of vehicle capacities, constraints (3.27) bound from below the cumulative quantity on the vehicle i at each container and enforce the vehicle capacity. Constraints (3.28) reset the cumulative quantity on the vehicle to zero at the origin, destination, and dumps. Keeping track of the cumulative quantity on the vehicle is achieved by constraints (3.29).

$$q_{ikt} \leq Q_{ikt} \leq \Omega_k, \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, i \in \mathcal{P} \quad (3.27)$$

$$Q_{ikt} = 0, \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, i \in \mathcal{N} \setminus \mathcal{P} \quad (3.28)$$

$$Q_{ikt} + q_{jkt} \leq Q_{jkt} + \Omega_k(1 - x_{ijkt}), \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, i \in \mathcal{N} \setminus \{d\}, j \in \mathcal{P} \quad (3.29)$$

The next four constraints express the intra-day temporal characteristics of the problem. Constraints (3.30) calculate the start-of-service time at each point. In addition, these constraints

eliminate the possibility of subtours and ensure that a point is not visited more than once by the same vehicle. Constraints (3.31) and (3.32) enforce the time windows. Constraints (3.33) provide a lower bound on the tour duration and apply the maximum tour duration.

$$S_{ikt} + \delta_i + \tau_{ijk} \leq S_{jkt} + (\mu_i + \delta_i + \tau_{ijk})(1 - x_{ijk}), \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, i \in \mathcal{N} \setminus \{d\}, j \in \mathcal{N} \setminus \{o\} \quad (3.30)$$

$$\lambda_i \sum_{j \in \mathcal{N}} x_{ijk} \leq S_{ikt}, \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, i \in \mathcal{N} \setminus \{d\} \quad (3.31)$$

$$S_{jkt} \leq \mu_j \sum_{i \in \mathcal{N}} x_{ijk}, \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, j \in \mathcal{N} \setminus \{o\} \quad (3.32)$$

$$0 \leq S_{dkt} - S_{okt} \leq H, \quad \forall t \in \mathcal{T}, k \in \mathcal{K} \quad (3.33)$$

Finally, constraints (3.34) and (3.35) establish the variable domains.

$$x_{ijk}, y_{ikt}, z_{kt} \in \{0, 1\}, \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, i, j \in \mathcal{N} \quad (3.34)$$

$$q_{ikt}, Q_{ikt}, I_{it}, S_{ikt} \geq 0, \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, i \in \mathcal{N} \quad (3.35)$$

3.2.3 Model Reformulations for Solving Benchmark Instances

The model presented above describes a problem that is not encountered in the literature. In order to evaluate the performance of the solution methodology presented in Section 3.3, we test it on IRP and VRP instances for problems with similar features. Below, we present the necessary reformulations of the original model, while the corresponding modifications to the general solution methodology are described in Section 3.3.3.

Reformulation for the Benchmark IRP.

For the benchmark IRP instances from the literature, we assume a distribution context and the presence of inventory holding costs at the depot and the customers. There are no intermediate facilities. We are in a deterministic setting and $\rho_{it} = \mathbb{E}(\rho_{it}), \forall t \in \mathcal{T}, i \in \mathcal{P}$. The commodity becomes available at the depot at a rate ρ_{ot} on day t and is consumed by customer i at a rate ρ_{it} on day t . Let η_o and η_i denote the inventory holding cost per day at the depot and customer i , respectively. In addition, we redefine q_{ikt} as the quantity delivered by vehicle k to customer i on day t , and Q_{ikt} as the cumulative quantity delivered by vehicle k arriving at point i on day t . The objective function of the benchmark IRP is composed of the inventory holding costs at the depot and the customers, and the routing cost, and writes as:

$$\begin{aligned} \min z^{\text{IRPB}} = & \sum_{t \in \mathcal{T} \cup \mathcal{T}^+} \eta_o I_{ot} + \sum_{t \in \mathcal{T} \cup \mathcal{T}^+} \sum_{i \in \mathcal{P}} \eta_i I_{it} \\ & + \sum_{t \in \mathcal{T}} \sum_{k \in \mathcal{K}} \left(\varphi_k z_{kt} + \beta_k \sum_{i \in \mathcal{N}} \sum_{j \in \mathcal{N}} \pi_{ij} x_{ijk} + \theta_k (S_{dkt} - S_{okt}) \right). \end{aligned} \quad (3.36)$$

A special set of constraints is needed for the inventory definition at the depot. Constraints (3.37) define the inventory level at the depot as the sum of the previous day's inventory and quantity made available minus the previous day's amount delivered to customers. Constraints (3.38) forbid a stock-out at the depot given the total quantity delivered to customers.

$$I_{0t} = I_{0(t-1)} + \rho_{0(t-1)} - \sum_{k \in \mathcal{K}} \sum_{i \in \mathcal{P}} q_{ik(t-1)}, \quad \forall t \in \mathcal{T}^+ \quad (3.37)$$

$$I_{0t} \geq \sum_{k \in \mathcal{K}} \sum_{i \in \mathcal{P}} q_{ikt}, \quad \forall t \in \mathcal{T} \quad (3.38)$$

We redefine the evolution of the inventory level at the customers for a distribution context. Constraints (3.39) define the inventory level at the customers as the sum of the previous day's inventory and quantity delivered minus the previous day's demand. Constraints (3.40) forbid the occurrence of stock-outs.

$$I_{it} = I_{i(t-1)} + \sum_{k \in \mathcal{K}} q_{ik(t-1)} - \rho_{i(t-1)}, \quad \forall i \in \mathcal{P}, t \in \mathcal{T}^+ \quad (3.39)$$

$$I_{it} \geq 0, \quad \forall i \in \mathcal{P}, t \in \mathcal{T}^+ \quad (3.40)$$

The order-up-to policy also needs to be redefined for a distribution context. Constraints (3.41), (3.42), and (3.43) express the fact that if a customer is visited, its inventory is brought up to its capacity.

$$q_{ikt} \geq \omega_i y_{ikt} - I_{it}, \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, i \in \mathcal{P} \quad (3.41)$$

$$q_{ikt} \leq \omega_i - I_{it}, \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, i \in \mathcal{P} \quad (3.42)$$

$$q_{ikt} \leq \omega_i y_{ikt}, \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, i \in \mathcal{P} \quad (3.43)$$

To avoid unnecessary complications in the reformulation, we can safely assume that there is a single dummy dump with zero service time and distance to the depot. The basic routing constraints (3.15)–(3.20), the vehicle capacity constraints (3.27)–(3.29), the intra-day temporal constraints (3.30)–(3.33), and the domain constraints (3.34)–(3.35) can thus be reused.

Reformulation for the Benchmark VRP.

For the VRP, it suffices to collapse the planning horizon to $\mathcal{T} = \{0\}$ and redefine the objective function z in terms of the Routing Cost (RC) only:

$$\min z^{\text{VRPB}} = \sum_{k \in \mathcal{K}} \left(\varphi_k z_{k0} + \beta_k \sum_{i \in \mathcal{N}} \sum_{j \in \mathcal{N}} \pi_{ij} x_{ijk0} + \theta_k (S_{dk0} - S_{ok0}) \right). \quad (3.44)$$

For day $t = 0$, demands are deterministic. As far as the constraints are concerned, the inequality sign in constraints (3.18) of the original model becomes an equality sign, providing

that each container should be visited by exactly one vehicle. Constraints (3.21) and (3.22) are dropped since the VRP is solved for a single period and we disregard its effect on the future. Constraints (3.23) are dropped as they become redundant given the modified constraints (3.18). Since there is no inventory tracking over the planning horizon, it is irrelevant whether we are in a collection or in a distribution context.

3.3 Adaptive Large Neighborhood Search

Adaptive Large Neighborhood Search (ALNS) was introduced by Ropke and Pisinger (2006a) in the context of the pickup and delivery problem with time windows. It is a type of large neighborhood search in which a number of operators compete in modifying the current solution. At each iteration of the search process, a number of customers is removed from the current solution by a destroy operator, after which they are reinserted elsewhere by a repair operator. In the context of our IRP, not all customers need to be visited every day, or even at all. Hence, we do not require that all removed customers be reinserted by the repair operator. The search guiding principle can be based on any meta-heuristic framework. Simulated annealing appears to be the preferred approach in the ALNS literature, and is also the one we implement. Given an incumbent solution s , a randomly drawn neighbor solution s' is always accepted if $f(s') < f(s)$, and with probability $\exp(-(f(s') - f(s))/T)$ otherwise, with $f(s)$ representing the solution cost and $T > 0$ the current temperature. The temperature is initialized as T^{start} and is reduced at each iteration by a cooling rate $r \in (0, 1)$. The search stops when T reaches a predetermined T^{end} .

Operator choice is governed by a roulette-wheel mechanism. Each operator i has a weight W_i , which depends on its past performance and a score. Given the set of destroy (repair) operators \mathcal{O} , the destroy (repair) operator i is selected with probability $W_i / \sum_{j \in \mathcal{O}} W_j$. The ALNS starts with all weights set to one and all scores set to zero. The scores of the selected destroy-repair couple are increased by e_1 if they find a new best feasible solution, by $e_2 < e_1$ if they improve the incumbent, and by $e_3 < e_2$ if they do not improve the incumbent but the new solution is accepted. This strategy rewards successful operator couples, while at the same time maintaining diversification during the search. It is important to note that if a destroy-repair couple leads to a visited solution, no reward is applied. The search is divided into segments of F iterations each, at the end of which the operator weights are updated. Let C_i^F denote the score of operator i and N_i^F the number of times it was applied in the last segment of length F . The new weights are computed as follows:

$$W_i = \begin{cases} W_i & \text{if } N_i^F = 0, \\ (1 - b)W_i + b C_i^F / (m_i N_i^F) & \text{otherwise.} \end{cases} \quad (3.45)$$

In expression (3.45), m_i is a normalization factor damping the weights of more computationally expensive operators by multiplying the number of times they were applied (Ropke and Pisinger, 2006b; Coelho et al., 2012a). The value $b \in [0, 1]$ is a reaction factor, controlling the

Algorithm 3.1: Adaptive Large Neighborhood Search

Input initial solution s^{init}
Output best found solution s^{best}

- 1: all weights equal to 1, all scores equal to 0
- 2: $s^{\text{best}} \leftarrow s \leftarrow s^{\text{init}}$
- 3: $T \leftarrow T^{\text{start}}$
- 4: **while** $T \geq T^{\text{end}}$ **do**
- 5: $s' \leftarrow s$
- 6: select a destroy-repair couple using roulette wheel and apply to s'
- 7: **if** $f(s') < f(s)$ **then**
- 8: $s \leftarrow s'$
- 9: **if** $f(s') < f(s^{\text{best}})$ **and** s' is feasible **then**
- 10: $s^{\text{best}} \leftarrow s'$
- 11: update scores of destroy-repair couple by e_1
- 12: **else**
- 13: update scores of destroy-repair couple by e_2
- 14: **end if**
- 15: **else if** s' is accepted
- 16: $s \leftarrow s'$
- 17: update scores of destroy-repair couple by e_3
- 18: **end if**
- 19: **if** iteration count is multiple of F **then**
- 20: update weights and reset scores to 0
- 21: **end if**
- 22: $T \leftarrow rT$
- 23: **end while**

relative effect of past performance and the scores on the new weights. Once the weights are updated, C_i^F and N_i^F are reset to zero. Algorithm 3.1 is a pseudocode of the ALNS implementation with simulated annealing. The function $f(\cdot)$ represents the full solution cost including penalties for feasibility violations, as explained in Section 3.3.1 next. Regarding the initial solution s^{init} , we build empty tours consisting of the depot as an origin and destination and one dump in between, without inserting any containers. An empty tour is built for each available vehicle on each day of the planning horizon. Since the destroy operators will have no effect in the beginning, the repair operators will insert containers and construct a non-empty solution.

3.3.1 Solution Representation

To facilitate the search and avoid becoming trapped in local optima, we admit infeasible intermediate solutions at a penalty. This relaxation technique is especially useful for tightly constrained problems. Let s be a solution and let $\mathcal{N}'_{k_t}(s)$ denote all point visits by vehicle k on day t in s , where each visit is a replication of the visited point. In addition, let $\mathcal{P}'_{k_t}(s) \subset$

$\mathcal{N}'_{kt}(s)$ denote all point visits where the next visit is a dump. We also define the function $(x)^+ = \max\{0, x\}$. Our ALNS admits the following types of intermediate feasibility violations:

1. Vehicle capacity violation is the sum of excess cumulative demand in $\mathcal{P}'_{kt}(s), \forall t \in \mathcal{T}, k \in \mathcal{K}$. Formally, it is defined as:

$$V^\Omega(s) = \sum_{t \in \mathcal{T}} \sum_{k \in \mathcal{K}} \sum_{i \in \mathcal{P}'_{kt}} (Q_{ikt} - \Omega_k)^+ \quad (3.46)$$

2. Time window violation is the total violation of the upper time window bounds μ_i in $\mathcal{N}'_{kt}(s), \forall t \in \mathcal{T}, k \in \mathcal{K}$. Lower time window bounds cannot be violated because if the vehicle arrives at point i before λ_i , it waits. Hence, formally, we have:

$$V^\mu(s) = \sum_{t \in \mathcal{T}} \sum_{k \in \mathcal{K}} \sum_{i \in \mathcal{N}'_{kt}} (S_{ikt} - \mu_i)^+ \quad (3.47)$$

3. Duration violation is expressed as the sum of excess durations. It is verified after time window violation. For each tour that has no time window violation, we apply forward time slack reduction (Savelsbergh, 1992), which may minimize tour duration while preserving time window feasibility (see Algorithm 2.1 in Chapter 2). In mathematical terms, duration violation writes as:

$$V^H(s) = \sum_{t \in \mathcal{T}} \sum_{k \in \mathcal{K}} (S_{dkt} - S_{okt} - H)^+ \quad (3.48)$$

4. Container capacity violation is the sum of excess container inventories $\forall t \in \mathcal{T}^+, i \in \mathcal{P}$, or:

$$V^\omega(s) = \sum_{t \in \mathcal{T}^+} \sum_{i \in \mathcal{P}} (I_{it} - \omega_i)^+ \quad (3.49)$$

5. Backorder limit violation is the sum of excess container inventories on day $t = 0, \forall i \in \mathcal{P}$ that are not visited on day $t = 0$. In mathematical terms, this is expressed as:

$$V^0(s) = \sum_{i \in \mathcal{P}} \left(\left(1 - \sum_{k \in \mathcal{K}} y_{ik0} \right) (I_{i0} - \omega_i)^+ \right). \quad (3.50)$$

6. Accessibility violation is the sum of inaccessible visits in $\mathcal{N}'_{kt}(s), \forall t \in \mathcal{T}, k \in \mathcal{K}$. They are accounted for as:

$$V^\alpha(s) = \sum_{t \in \mathcal{T}} \sum_{k \in \mathcal{K}} \sum_{i \in \mathcal{N}'_{kt}} (y_{ikt} - \alpha_{ik})^+ \quad (3.51)$$

Including the possibility of all violations, the complete solution cost during the search is

represented by:

$$f(s) = z(s) + L^\Omega V^\Omega(s) + L^\mu V^\mu + L^H V^H(s) + L^\omega V^\omega(s) + L^0 V^0(s) + L^\alpha V^\alpha(s). \quad (3.52)$$

The parameters L^Ω through L^α are the penalties for each type of feasibility violation. They are dynamically adjusted during the search so as to encourage the exploration of infeasible solutions but to avoid staying infeasible for too long. At each accepted solution, the incumbent s is checked for each type of violation. If it is non-zero, its respective penalty is multiplied by a rate $\ell > 1$, otherwise it is divided by the same rate. If s has no feasibility violation, the values of $f(s)$ and $z(s)$ coincide. As indicated in Algorithm 3.1, the update of the best solution requires feasibility with respect to conditions (3.46) through (3.51).

3.3.2 Operators

The main ingredient of the ALNS are the destroy and repair operators. Some of the operators are inspired or adapted from the literature (Ropke and Pisinger, 2006a,b; Coelho et al., 2012a; Buhrkal et al., 2012; Hemmelmayr et al., 2013), while others are developed to capture the specifics of our problem, in particular the stochastic objective function and the presence of a heterogeneous fixed fleet. The choice of operators balances between diversification and intensification. We use the following destroy operators:

1. *Remove ν containers randomly.* This operator selects a random tour and removes a random container from it. It is applied ν times, where ν is an integer drawn from a discrete semi-triangular distribution bounded below by 1 and above by the number of containers in \mathcal{P} . Small ν 's result in cosmetic changes to the solution, while big ν 's, which are drawn with a lower probability, lead to larger perturbations.
2. *Remove ν worst containers.* This operator removes the container that would lead to the largest savings Δf^{\max} in the solution cost. It is applied ν times.
3. *Shaw removals with relatedness.* Based on the ideas of Shaw (1997) and Ropke and Pisinger (2006a), this operator removes containers based on a relatedness measure among them. It starts by selecting a random tour and a random container i in this tour, and computing the relatedness R_{ij} of i to all containers j in the tours scheduled on the same day t as the randomly selected tour. We define the relatedness measure R_{ij} of container i to container j as:

$$R_{ij} = d_1 \pi_{ij}^{[0,1]} + d_2 (|\lambda_i - \lambda_j| + |\mu_i - \mu_j|)^{[0,1]} + d_3 |o_{it} - o_{jt}|^{[0,1]}, \quad (3.53)$$

where the first term captures the distance, the second terms captures the time window difference and the third term captures the overflow probability difference on day t . As in expression (3.11) for the EOEC, the latter is given by:

$$o_{it} = \mathbb{P}(\sigma_{it} = 1 \mid m = \max(0, g \in \mathcal{T} : g < t : \exists k \in \mathcal{K} : y_{ikg} = 1)). \quad (3.54)$$

These terms are scaled between zero and one, as indicated in superscript, and weighted by the parameters d_1 , d_2 and d_3 . The relatedness measures are again scaled between zero and one. Container i and all containers j for which R_{ij} is less than a threshold d_4 are removed.


4. *Remove container cluster.* Inspired by the work of Ropke and Pisinger (2006b), this operator removes large clusters of containers. It selects a random day t in the planning horizon and divides the containers visited on this day into k clusters, where k is chosen to be the number of tours executed on this day. If there is only one tour, its containers are divided into 2 clusters. Clustering is performed using Kruskal's algorithm, which progressively merges the containers into clusters based on distance, until the required number of clusters k is reached. Finally, a cluster is chosen randomly and removed as long as it contains less than half of the containers visited on day t .
5. *Empty a random day.* This operator selects a random day and empties all tours performed on it.
6. *Empty a random vehicle.* This operator selects a random vehicle and empties the tours performed by it on all days.
7. *Remove a random dump.* This operator selects a random tour and a random dump in it, excluding the last dump, and removes it.
8. *Remove the worst dump.* This operator removes the dump that would lead to the largest savings Δf^{\max} in the solution cost. The last dump in each tour is never removed.
9. *Remove consecutive visits.* This operator inspects each container over the planning horizon and, if it is visited on two consecutive days, removes the second visit. This is based on the idea that optimal or good-quality solutions will rarely visit the same container on consecutive days.

In addition, we use the following repair operators:

1. *Insert ν containers randomly.* This operator selects a random tour and a random container from \mathcal{P} not visited on the day the tour is performed, and inserts it in the tour using best insertion, i.e. in the position in the selected tour that would lead to the minimum increase Δf^{\min} in the solution cost. It is applied ν times.
2. *Insert ν containers in the best way.* This operator identifies for each container $i \in \mathcal{P}$ the tour and the position in that tour that would lead to the minimum increase Δf_i^{\min} in the solution cost if the container is inserted there, checking that the container is not visited on the day the tour is performed. The containers in \mathcal{P} are sorted in ascending order of Δf_i^{\min} and the first ν of them are inserted in the previously identified tours and positions.

3. *Insert ν containers with regret- k .* As noted in Ropke and Pisinger (2006a), the motivation for using regret is to introduce a look-ahead information in the insertion process. Let Y_{ik} indicate the tour in which inserting container i using best insertion leads to the k^{th} lowest increase in the solution cost $\Delta f_{i, Y_{ik}}$. For a container i , we define the regret- k value as $c_i^k = \Delta f_{i, Y_{ik}} - \Delta f_{i, Y_{i1}}$, i.e. the difference between inserting the container in its best tour and its k^{th} best tour. It may be impossible to insert some containers in k different tours, thus the regret is computed for the largest possible $k' \leq k$. The containers in \mathcal{P} are sorted in ascending order of k' and descending order of $c_i^{k'}$. The first ν containers in the ordered list are inserted in the tours and positions that would lead to the minimum increase Δf^{\min} in the solution cost. In other words, we insert the containers that we will regret the most if they are not inserted now.
4. *Shaw insertions with relatedness.* This operator selects a random day t and a random container $i \in \mathcal{P}$ not visited on day t . It then proceeds to find the relatedness measure R_{ij} , as defined by formula (3.53), to all containers $j \in \mathcal{P}$ also not visited on day t . It inserts the container i as well as all containers j not visited on day t , for which R_{ij} is lower than a threshold d_A , in the tours executed on day t and in the positions that would lead to the minimum increase Δf^{\min} in the solution cost.
5. *Swap ν random containers.* This operator selects two random tours and a random container in each one, and swaps the container-to-tour assignment by using best insertion in each tour. It is applied ν times.
6. *Insert a dump randomly.* This operator selects a random tour and a random dump from \mathcal{D} and inserts it at a random position in the tour.
7. *Insert a dump in the best way.* This operator selects a random dump from \mathcal{D} and inserts it in the tour and in the position that would lead to the minimum increase Δf^{\min} in the solution cost.
8. *Swap random dumps.* This operator selects two random tours and a random dump in each one, and swaps the dumps.
9. *Replace a random dump.* This operator selects a random tour and a random dump in it, and replaces it with another random dump from \mathcal{D} .
10. *Reorder dumps.* Based on the idea of Hemmelmayr et al. (2013), this operator selects a random tour, removes all dump visits from it, and finds the locally optimal dump visit configuration that preserves vehicle capacity feasibility. Figure 3.2 provides an illustrative example of a tour starting at the depot, visiting containers i_1 through i_5 , and terminating at the depot. The values of ρ_1 through ρ_5 denote the container demands, and we assume a vehicle with a capacity of 10 units. Because a dump will never be visited between the depot and the first container, they can be merged into a single node. Each arc starts at a container and ends at a container or the depot, visiting on its way the indicated containers and the best dump, either j_1 or j_2 , before the end

Figure 3.2: Feasibility Graph of the Reorder Dumps Operator



node. The resulting directed graph is not necessarily complete, as it only contains the vehicle capacity preserving arcs. The solution to the problem amounts to finding the shortest path from the origin to the destination node representing the depot. We use the Bellman-Ford algorithm and post-optimize the result using 2-opt local search.

The destroy operators that empty a random day and a random vehicle leave the affected tour with the depot as an origin and destination, and a dump, and the cost of such a tour is considered zero. Thus, all original tours always remain available during the search for removal of points from or insertion of points into. This is a straightforward way to manage the presence of a heterogeneous fixed fleet without having to re-evaluate periodically vehicle-to-tour assignments. This strategy will likely not be applicable to more classical meta-heuristics that exploit much smaller neighborhoods.

3.3.3 Algorithmic Modifications for Solving Benchmark Instances

Several modifications to the original ALNS algorithm are needed in order to integrate the model reformulations described in Section 3.2.3 and necessary for solving the benchmark IRP and VRP instances.

Modifications for the Benchmark IRP.

The benchmark IRP considers no intermediate facilities and so we disregard all dump-related operators. To avoid further changes to the algorithm, the always-present dump visit before the destination depot is created as a dummy node with zero service time and distance to the depot, and as such does not affect the solution. For a given solution s , we define the depot inventory violation as expressed by constraints (3.38) in the benchmark IRP reformulation:

$$V^I(s) = \sum_{t \in \mathcal{T}} \left(\sum_{k \in \mathcal{K}} \sum_{i \in \mathcal{P}} q_{ikt} - I_{ot} \right)^+ \quad (3.55)$$

The violation $V^I(s)$ is multiplied by parameter L^I and added to the objective function representation $f(s)$ as in expression (3.52). Additionally, we redefine the container capacity

violation (3.49) in terms of stock-out as it applies to a distribution context:

$$V^\omega(s) = \sum_{t \in \mathcal{T}^+} \sum_{i \in \mathcal{P}} (-I_{it})^+ \quad (3.56)$$

The back-order violation (3.50) is dropped since back-orders do not apply to the benchmark IRP.

Modifications for the Benchmark VRP

In the original ALNS algorithm, the number of containers inserted into the solution by a repair operator is randomly drawn and not necessarily the same as the number of containers removed by the destroy operator applied before it. This design allows to vary the number of containers visited each day, as this is a decision variable in the IRP. Contrarily, the VRP assumes that all containers are visited in the solution. To achieve the latter, we implement an initial solution construction procedure and a simple rearrangement of the destroy and repair operators.

To construct an initial solution, we use repair operator number 1, *insert ν containers randomly*, to insert all containers into the solution. The resulting initial solution is not necessarily feasible. Then we redefine the operators so that all destroy operators and repair operators 5 through 10 are now drawn first, while repair operators 1 through 4 are drawn second. This separation is based on the operators' ability to reinsert containers into the solution. In other words, the repair operators are now only those that have this ability, namely *insert ν containers randomly*, *insert ν containers in the best way*, *insert ν containers with regret- k* , and *Shaw insertions with relatedness*. Moreover, the number of containers to be reinserted is not random. The repair operators now reinsert all containers that were previously removed by the destroy operator. If the destroy operator did not remove any containers, the repair operator is not applied. Given that all containers are now visited in the solution, we drop violations (3.49) and (3.50) from the solution representation, i.e. container capacity violation and backorder violation. If the problem at hand considers no intermediate facilities, we disregard all dump-related operators. The always-present dump visit before the destination depot in this latter case is created as a dummy node with zero service time and distance to the depot, and as such does not affect the solution.

3.4 Numerical Experiments

The ALNS is implemented as a single-thread application in Java, and the forecasting model and the probability calculator for the state probability tree (Figure 3.1) are scripted in R. All tests have been carried out on a 3.33 GHz Intel Xeon X5680 server running a 64-bit Ubuntu 16.04.2. Each instance is solved 10 times, out of which we report the best, average, and worst result, or averaged values of the best, average and worst result over a set of instances, unless

indicated otherwise. Section 3.4.1 describes how the algorithmic parameters were tuned. This is followed by results of the ALNS performance on benchmark IRP and VRP instances in Section 3.4.2. Finally, Section 3.4.3 presents an extensive analysis of the model and solution methodology applied on a case study with instances derived from real data.

3.4.1 Parameter Tuning

The algorithmic parameters were tuned on the Archetti et al. (2007) instances and the case study instances described below. We first tuned the SA-related parameters followed by the ALNS-related and the operator-related parameters. Initial values were either borrowed from ALNS implementations in the literature or based on preliminary trial-and-error combinations. The parameters were tuned one by one, unless indicated otherwise, in the order in which they appear in Table 3.2. The initial temperature was set sufficiently high for an initial feasible solution to be found without difficulty. Once this is the case, the temperature is calibrated so that the probability of accepting a solution which is worse than it by a factor of w is 50%. The purpose of this strategy is to limit the search at very high temperatures (Ropke and Pisinger, 2006a). The cooling rate typically results in several hundred thousand iterations on the Archetti et al. (2007) instances, and the final temperature allows sufficient time for the algorithm to converge. The penalty change rate multiplies or divides the penalties associated with conditions (3.46) through (3.51) as explained in Section 3.3.1. After fixing the SA-related parameters, we tuned the ALNS-related parameters. The rewards were tuned together, and after testing several configurations we chose one that attributes a relatively lower reward e_3 for a non-improving but accepted solution. The two destroy operators *Shaw removals with relatedness* and *remove container cluster* were given normalization factors m_i of 8, and the two repair operators *insert v containers in the best way* and *insert v containers with regret- k* were given normalization factors m_i of 4.5. The normalization factors for the rest are all equal to one. For the operator-related parameters, the best results were obtained for regret-2. The relatedness weights d_1 , d_2 and d_3 were calibrated at 0.54, 0.23 and 0.23, respectively. The relatedness thresholds of 0.2 for removals and 0.3 for insertions were found to perform the best.

Table 3.2: Algorithmic Parameters

SA-Related		ALNS-Related		Operator-Related	
Parameter	Value	Parameter	Value	Parameter	Value
Initial temperature (T^{start})	10,000	F segment length	2000	Rel. weight d_1	0.54
Start temp. control param. (w)	0.6	Reaction factor (b)	0.5	Rel. weight d_2	0.23
Cooling rate (r)	0.99998	Reward e_1	30	Rel. weight d_3	0.23
Final temperature (T^{end})	0.01	Reward e_2	20	Rel. threshold d_4	0.2/0.3
Penalty change rate (ℓ)	1.06	Reward e_3	5	Regret k	2

3.4.2 Benchmark Results

The model reformulations and algorithmic modifications presented in Section 3.2.3 and Section 3.3.3, respectively, were necessary in order to assess the performance of the solution methodology on benchmark instances from the literature. Below we present the results obtained by the ALNS on classical IRP and VRP instance testbeds. Archetti et al.'s (2007) instances test the ALNS in a classical IRP context with order-up-to level policy. Crevier et al.'s (2007) instances test the ALNS in a VRP setting with intermediate facilities, the latter being an important feature that is present in our problem in the form of dumps. Crevier et al.'s (2007) instances are also representative of a standard waste collection problem and used by Hemmelmayr et al. (2013) in this context. Finally, Taillard's (1999) instances test the ALNS in a VRP context with a rich heterogeneous fixed fleet, which is the most general type of fleet. These benchmark instances capture many of the features of our problem, albeit not simultaneously.

Results on IRP Benchmarks.

We ran the ALNS algorithm on the IRP benchmark set proposed by Archetti et al. (2007), which is the first classical IRP testbed in the literature. It represents a deterministic IRP in a distribution context where an inventory holding cost η_o applies at the depot, and inventory holding costs η_i apply at the customers. There is a single vehicle available each day, with its daily deployment cost φ_k and unit-time running cost θ_k both equal to zero, and its unit-distance running cost $\beta_k = 1$. Stock-outs are forbidden at the customers and the depot. Vehicle tours are only limited by the vehicle capacity, and no rich VRP features such as intermediate facilities, time windows, or a maximum tour duration are considered.

The set includes two equal subsets with high, respectively low, inventory holding costs η_i . The length of the planning horizon $|\mathcal{T}|$ is either 3 or 6 periods, and the number of customers n varies from 5 to 50 for $|\mathcal{T}| = 3$, and from 5 to 30 for $|\mathcal{T}| = 6$. Five instances are generated for each combination of η_i , $|\mathcal{T}|$ and n , thus resulting in a total of 160 instances. Using a branch-and-cut algorithm, Archetti et al. (2007) solve with a proof of optimality all instances except one (low η_i , $|\mathcal{T}| = 3$, $n = 50$), where the gap is brought to 0.99% within the time limit of two hours. A number of heuristic algorithms are tested on these instances or derivations thereof (Archetti et al., 2012; Coelho et al., 2012a,b). The most successful one is the hybrid heuristic of Archetti et al. (2012) which is able to achieve an optimality gap of 0.1% for the order-up-to policy based on a single experiment per instance, and with computation times up to several thousand seconds for the largest instances on an Intel Dual Core 1.86 GHz processor.

Table 3.3 presents our results on the Archetti et al. (2007) instances. The first two columns report the number of periods $|\mathcal{T}|$ in the planning horizon and the number of customers n . The remainder of the table is divided into two parts, with results for the instances with high, respectively low, inventory holding cost, each providing the average runtime in seconds, and

Table 3.3: Results on Archetti et al. (2007) Instances

$ \mathcal{T} $	n	High Inventory Holding Cost				Low Inventory Holding Cost			
		Runtime(s.)	Best Gap(%)	Avg Gap(%)	Worst Gap(%)	Runtime(s.)	Best Gap(%)	Avg Gap(%)	Worst Gap(%)
3	5	69.08	0.00	0.00	0.00	85.69	0.00	0.00	0.00
3	10	183.94	0.00	0.00	0.00	156.36	0.00	0.00	0.00
3	15	317.93	0.00	0.00	0.00	274.05	0.00	0.00	0.00
3	20	440.02	0.00	0.00	0.01	444.68	0.00	0.00	0.02
3	25	523.42	0.00	0.08	0.25	501.78	0.01	0.20	0.66
3	30	835.21	0.01	0.15	0.32	649.09	0.00	0.41	0.98
3	35	866.06	0.00	0.15	0.36	731.21	0.00	0.46	1.68
3	40	896.91	0.02	0.18	0.44	976.83	0.16	0.47	0.97
3	45	1124.57	0.05	0.42	0.91	1074.19	0.00	1.05	2.53
3	50	1424.27	0.06	0.35	0.79	1223.56	0.13	1.19	2.15
6	5	105.86	0.00	0.00	0.00	73.28	0.00	0.00	0.00
6	10	184.48	0.00	0.01	0.08	181.93	0.00	0.00	0.00
6	15	333.82	0.01	0.09	0.15	272.03	0.00	0.03	0.16
6	20	394.39	0.00	0.17	0.41	420.28	0.05	0.34	0.82
6	25	636.27	0.12	0.34	0.82	546.85	0.09	0.67	1.60
6	30	725.63	0.10	0.47	0.93	733.12	0.44	1.43	2.63
Average		566.37	0.02	0.15	0.34	521.56	0.05	0.39	0.89

the best, average and worst gap to the optimal solution obtained over 10 runs. Each row averages the latter over the five instances for each combination of η_i , $|\mathcal{T}|$ and n . Our results are comparable to the best from the literature. The ALNS attains a best gap of 0.02% and 0.05%, an average gap of 0.15% and 0.39%, and a worst gap of 0.34% and 0.89% on the high and low inventory holding cost instances, respectively. In comparison, Archetti et al.’s (2012) algorithm obtains a gap of 0.06% and 0.10%, respectively, for a single run per instance. We are able to solve to optimality almost all instances with up to 35 customers for $|\mathcal{T}| = 3$, and with up to 15 customers for $|\mathcal{T}| = 6$. Similar quality results can also be found in Coelho et al. (2012a) and Coelho et al. (2012b), also when it comes to the higher gaps on the low inventory holding cost instances. A possible explanation could be that for low inventory holding costs the importance of the container selection decision in each period becomes relatively more pronounced. Our computation times are also very competitive compared to those in the literature, although a more rigorous scaling approach could be difficult due to the lack of precise processor architecture specifications in some of the works.

Results on VRP Benchmarks.

The SIRP that we study includes a rich routing component. Since the routing component in the IRP benchmarks under consideration is very simple, we test our ALNS on two VRP benchmark instance sets, namely those of Crevier et al. (2007) and Taillard (1999).

Crevier et al. (2007) solve the Multi-Depot VRP with Inter-depot routes (MDVRPI). Their instances consist of two sets of randomly generated instances with intermediate facilities, a homogeneous fixed fleet, and a maximum tour duration. Each vehicle’s daily deployment cost φ_k and unit-time running cost θ_k are both equal to zero, and its unit-distance running cost $\beta_k = 1$. The set (a1–l1) includes 12 newly generated instances with two to five interme-

3.4. Numerical Experiments

mediate facilities and 48 to 216 customers. The set (a2–j2) includes 10 instances derived from those of Cordeau et al. (1997) by adding a central depot where the vehicles are stationed. It contains four to six intermediate facilities and 48 to 288 customers. The Best Known Solutions (BKS) to both sets are obtained by Hemmelmayr et al. (2013) who use a VNS with the dynamic programming procedure for the insertion of the intermediate facilities presented in Section 3.3.2. In Table 3.4, the instance name is followed by the number of customers n , the computation time in seconds, the best and average cost obtained by Hemmelmayr et al. (2013) over 10 runs. The next three columns report the values produced by our ALNS. The last two columns represent the percent gap of our best and our average cost with respect to those of Hemmelmayr et al. (2013). Our best results are on average 0.40% from those of Hemmelmayr et al. (2013) and we are able to reach five of the BKS. Our gap with respect to the average value over 10 runs is under 1%. We note that the obtained values are better than those of the MNS in Chapter 2. We also observe that our ALNS is about six times slower than Hemmelmayr et al.’s (2013) VNS on these instances.

Taillard (1999) formalizes the Heterogeneous Fixed Fleet VRP (HFFVRP). The version we solve, known as the HFFVRP with fixed and variable costs, considers a non-zero daily deployment cost φ_k and unit-distance running cost β_k , and a zero unit-time running cost θ_k . The instance set is derived from the eight largest Golden et al. (1984) instances by specifying φ_k and β_k for each vehicle k so that no single vehicle is better than any other in terms of its capacity to cost ratio. The instances include 50, 75, and 100 customers, three to six vehicle types and up to six vehicles per type. Taillard (1999) spurred a strong scientific

Table 3.4: Results on Crevier et al. (2007) Instances

Instance	n	Hemmelmayr et al. (2013)			ALNS				
		Runtime(s.)	Best Cost	Avg Cost	Runtime(s.)	Best Cost	Avg Cost	Gap Best(%)	Gap Avg(%)
a2	48	73.80	997.94	997.94	140.41	997.94	998.17	0.00	0.02
b2	96	384.60	1291.19	1291.19	681.63	1291.19	1293.11	0.00	0.15
c2	144	900.60	1715.60	1715.84	2314.81	1733.95	1740.11	1.07	1.41
d2	192	1808.40	1856.84	1860.92	9264.81	1870.99	1879.90	0.76	1.02
e2	240	2958.60	1919.38	1922.81	17,694.38	1929.05	1943.05	0.50	1.05
f2	288	4274.40	2230.32	2233.43	32,170.43	2247.60	2273.39	0.77	1.79
g2	72	222.60	1152.92	1153.17	475.02	1152.92	1153.15	0.00	0.00
h2	144	939.60	1575.28	1575.28	2496.57	1582.78	1587.06	0.48	0.75
i2	216	2515.20	1919.74	1922.24	13,896.84	1947.78	1964.21	1.46	2.18
j2	288	4402.80	2247.70	2250.21	40,936.64	2259.25	2283.23	0.51	1.47
a1	48	85.20	1179.79	1180.57	200.24	1190.12	1197.22	0.88	1.41
b1	96	383.40	1217.07	1217.07	925.46	1218.09	1218.97	0.08	0.16
c1	192	1224.00	1866.76	1867.96	6435.66	1874.47	1884.54	0.41	0.89
d1	48	94.20	1059.43	1059.43	139.35	1059.43	1061.86	0.00	0.23
e1	96	373.20	1309.12	1309.12	743.72	1309.12	1319.97	0.00	0.83
f1	192	1536.00	1570.41	1573.05	7303.54	1571.62	1587.91	0.08	0.94
g1	72	202.80	1181.13	1183.32	308.81	1185.99	1189.76	0.41	0.54
h1	144	876.60	1545.50	1548.61	2622.40	1553.12	1565.19	0.49	1.07
i1	216	2014.80	1922.18	1923.52	10,383.43	1925.19	1935.39	0.16	0.62
j1	72	166.80	1115.78	1115.78	489.83	1116.67	1123.14	0.08	0.66
k1	144	873.60	1576.36	1577.96	2937.01	1580.84	1591.56	0.28	0.86
l1	216	2128.80	1863.28	1869.70	9343.72	1870.87	1884.84	0.41	0.81
Average		1292.73	1559.71	1561.32	7359.30	1566.77	1576.17	0.40	0.86

Table 3.5: Results on Taillard (1999) Instances

Instance	n	BKS	ALNS						
			Runtime(s.)	Best Cost	Avg Cost	Worst Cost	Best Gap(%)	Avg Gap(%)	Worst Gap(%)
13	50	3185.09	317.91	3228.97	3250.87	3279.00	1.38	2.07	2.95
14	50	10,107.53	335.92	10,123.50	10,143.21	10,158.54	0.16	0.35	0.50
15	50	3065.29	318.48	3072.59	3082.41	3090.24	0.24	0.56	0.81
16	50	3265.41	326.63	3292.22	3312.42	3322.43	0.82	1.44	1.75
17	75	2076.96	762.52	2097.29	2132.52	2148.15	0.98	2.68	3.43
18	75	3743.58	847.68	3801.22	3832.18	3860.08	1.54	2.37	3.11
19	100	10,420.34	1950.83	10,478.69	10,500.96	10,515.97	0.56	0.77	0.92
20	100	4761.26	1598.25	4901.45	4927.19	4946.61	2.94	3.49	3.89
Average		5078.18	807.27	5124.49	5147.72	5165.13	1.08	1.71	2.17

interest in this problem resulting in at least a dozen algorithms in the literature. The proof of optimality of the solutions to the 50- and 75-customer instances of the problem is due to Baldacci and Mingozzi (2009). In Table 3.5, the instance name is followed by the number of customers n and the BKS, which are due to multiple authors. Next are the computation time in seconds, the best, average and worst cost obtained by our ALNS. The last three columns report the percent gap of our best, average and worst cost with respect to the BKS. Our results are on average in the order of 2% from the BKS, most of which are proved to be optimal. Computation times are in the order of five to 30 minutes.

3.4.3 Case Study

In this section, we analyze the performance of our ALNS on sets of stochastic IRP instances derived from real data. In the text below, we first introduce the experimental setup and describes the instances. Then, we evaluate the effect of including the probability information in the objective function in terms of its impact on the expected cost and the frequency of occurrence of container overflows and route failures. Next, we compare the probabilistic approach to alternative practical deterministic policies such as artificial buffer capacities at the containers and trucks. In the two latter cases, we use simulation of the stochastic demands to test the quality of the solution produced by the ALNS. Finally, we employ a daily rolling horizon approach and derive empirical lower and upper bounds on the resulting cost over a given planning horizon.

Instances.

The case study data, except for the rolling horizon approach, includes 63 instances of white glass collections performed by a specific collection firm in the canton of Geneva, Switzerland. A map of the collection area was presented in Chapter 2. The instances are created using the historical records for weekly visits to white glass containers for the years 2014, 2015 and 2016. The planning horizon is seven days long, starting on Monday and finishing on Sunday. As established by constraints (3.22) in the mathematical model, there should be

no expected overflows on the first day after the end of the planning horizon, in our case the following Monday. On average, there are 41 containers per instance, and the maximum is 53, and their volumes range from 1000 to 3000 liters. Collection takes three or five minutes depending on container type. There are two dumps located far apart in the periphery of the city of Geneva. The fleet consists, depending on the instance, of one or two heterogeneous vehicles of volume capacity in the order of 30,000 liters and weight capacity of 10,000 to 15,000 kg, which are not available on the weekend. We also have access to historical waste levels for each container. Thus, the demands for each instance are forecast by the model from Section 3.2.1 using the previous 90 days of observations for each container. Two deposit sizes—two and ten liters—are used. For each instance, there is a distinct forecasting error ζ estimated by formula (3.4). We do not have information about tour duration, time windows and the cost parameters, for which we set realistic or reasonable values. Thus, tours should respect a maximum duration of four hours each, and the time windows correspond to 8:00 a.m. until noon. For the trucks, we use a daily deployment cost of 100 CHF, a cost of 2.95 CHF per kilometer and a cost of 40 CHF per hour. The overflow cost, which is normally determined by the municipality, is set to 100 CHF.

Probabilistic policies.

To study the impact of the probability information included in the objective function, here we perform two types of experiments on the instances described above. The first type considers the complete objective function with all relevant costs, as defined by expression (3.9). We label the problem with this objective "Complete". The second type minimizes the routing cost defined by expression (3.10), ignoring all costs related to container overflows, emergency collections and route failures, and we label the problem with the latter objective "Routing-only". Since the routing-only problem ignores all stochastic information and only the stochastic information, it becomes the deterministic version of the stochastic problem. Tables 3.6, 3.7 and 3.8 summarize the numerical results for various choices of the Emergency Collection Cost (ECC) and the Route Failure Cost Multiplier (RFCM), and each row represents averaged values over the 63 instances. In these three tables, the first three columns identify the type of objective considered (complete vs. routing-only), the applied ECC, and the applied RFCM. In Table 3.6, the next four columns report the computation time, the average number of tours, container collections and dump visits. As each instance is solved 10 times, the next three columns report the average over the 63 instances of the best, average and worst cost, respectively, over the 10 runs for each instance. The last two columns show the percent gap between the average and the best cost, and the worst and the best cost, respectively. We observe that computation times are in the order of 10 to 15 minutes, which is acceptable for an operational problem that is solved on a daily basis. The results indicate clearly that the complete objective solution collects on average more than twice as many containers and, as a consequence, performs more tours and dump visits. In terms of expected cost, it is 50 to 60% more expensive. Since the optimal solution is not available, we can only judge the quality of the result based on the gaps presented in the last two columns. Over the 63

Chapter 3. The Waste Collection IRP with Stochastic Demands

instances, the average gap between the average cost and the best cost is in the order of 0.5%, and between the worst cost and the best cost it is in the order of 1.5%, which is an indicator of the quality of our ALNS. The values are lowest for the routing-only objective and grow with higher ECC for the complete objective, reflecting the more challenging search space produced by the non-linearities present there. On the other hand, it appears that the gaps are almost unaffected by the RFCM.

Table 3.7 is a more detailed breakdown of the cost and efficiency structure of the set of objective functions presented in Table 3.6. The fourth, fifth and sixth column decompose the average solution cost from Table 3.6 into routing, overflow and route failure cost. The last three columns report the total collected volume in liters, and the volume per unit of total cost and per unit of routing cost, which can be regarded as performance indicators. The results reveal that the routing cost of the complete objective solution is on average only 30 to 40% higher than that of the routing-only objective solution. The rest of the difference in the total solution cost is explained by the contribution of the expected overflow cost. The routing cost is lower for a lower emergency collection cost, while the expected overflow cost remains almost unchanged. A higher emergency collection cost necessitates more frequent

Table 3.6: Probabilistic Policies: Basic Results for Cost Analysis on Real Data Instances

Objective	ECC	RFCM	Runtime(s.)	Avg Num Tours	Avg Num Collections	Avg Num Dump Visits	Best Cost (CHF)	Avg Cost (CHF)	Worst Cost (CHF)	Gap Avg-Best(%)	Gap Worst-Best(%)
Complete	100.00	1.00	870.65	1.95	44.41	2.24	662.65	666.64	672.87	0.60	1.54
Complete	100.00	0.50	871.84	1.95	44.45	2.25	662.38	666.57	673.30	0.63	1.65
Complete	100.00	0.25	885.52	1.95	44.46	2.24	662.38	666.92	673.15	0.69	1.63
Complete	100.00	0.00	871.81	1.95	44.46	2.23	662.26	666.78	674.01	0.68	1.78
Complete	50.00	1.00	864.57	1.92	42.39	2.18	648.14	651.36	656.77	0.50	1.33
Complete	50.00	0.50	855.51	1.92	42.40	2.17	647.99	651.50	656.90	0.54	1.37
Complete	50.00	0.25	873.28	1.92	42.36	2.16	648.05	651.15	656.66	0.48	1.33
Complete	50.00	0.00	856.39	1.92	42.35	2.18	648.14	651.40	656.47	0.50	1.29
Complete	25.00	1.00	841.94	1.90	41.03	2.16	638.61	641.41	646.06	0.44	1.17
Complete	25.00	0.50	844.22	1.90	41.05	2.16	638.38	641.22	645.89	0.44	1.18
Complete	25.00	0.25	846.67	1.90	41.01	2.15	638.57	641.50	646.19	0.46	1.19
Complete	25.00	0.00	855.83	1.90	41.01	2.15	638.42	641.49	646.36	0.48	1.24
Routing-only	0.00	0.00	681.27	1.83	16.64	1.87	421.99	422.48	423.12	0.12	0.27

Table 3.7: Probabilistic Policies: Key Performance Indicators for Cost Analysis on Real Data Instances

Objective	ECC	RFCM	Avg Routing Cost (CHF)	Avg Overflow Cost (CHF)	Avg Rte Failure Cost (CHF)	Avg Collected Volume (L)	Liters Per Unit Cost	Liters Per Unit Routing Cost
Complete	100.00	1.00	579.75	86.86	0.03	47,821.12	71.73	82.49
Complete	100.00	0.50	579.84	86.65	0.07	47,920.02	71.89	82.64
Complete	100.00	0.25	580.16	86.71	0.04	47,925.52	71.86	82.61
Complete	100.00	0.00	579.93	86.85	0.00	47,872.93	71.80	82.55
Complete	50.00	1.00	563.52	87.83	0.01	46,247.51	71.00	82.07
Complete	50.00	0.50	563.03	88.40	0.08	46,327.89	71.11	82.28
Complete	50.00	0.25	562.19	88.91	0.05	46,380.87	71.23	82.50
Complete	50.00	0.00	563.34	88.06	0.00	46,404.74	71.24	82.37
Complete	25.00	1.00	553.80	87.59	0.02	45,215.18	70.49	81.64
Complete	25.00	0.50	553.74	87.42	0.07	45,279.65	70.62	81.77
Complete	25.00	0.25	553.77	87.68	0.06	45,281.71	70.59	81.77
Complete	25.00	0.00	553.53	87.96	0.00	45,347.30	70.69	81.92
Routing-only	0.00	0.00	422.48	0.00	0.00	24,955.14	59.07	59.07

3.4. Numerical Experiments

visits as an attempt to further limit overflows. The route failure cost in both solutions is practically null. Not surprisingly, the solutions with the complete objective collect more volume as well. However, a better indication of their efficiency is provided by the collected volume per unit cost, which is 20% higher with respect to the total cost, and almost 40% higher with respect to the routing cost.

The relevance of the probability information captured by the objective function can be evaluated through the analysis of the occurrence of extreme events. After solving each instance, we simulate 10,000 scenarios, sampling the forecasting error independently for each container and each day using the estimate ζ . We then evaluate the effect on the occurrence of container overflows and route failures in the solution of the ALNS algorithm. An overflow is counted on each day, i.e. if a container is overflowing on two consecutive days because it is not collected, we count two overflows. Table 3.8 summarizes the number of overflows and route failures at the 75th, 90th, 95th and 99th percentiles of the 10,000 simulated scenarios, where each row is an averaged result for the 63 instances. We observe a strong negative correlation of the average number of overflows with the emergency collection cost and of the average number of route failures with the RFCM. What is more striking, however, is the difference between the series of complete objectives on the one hand and the routing-only objective on the other. While the complete objectives limit the number of overflows to about four, even at the extreme of the simulated scenarios, the average number of overflows for the routing-only objective is higher by a degree of magnitude.

Figure 3.3 is a visual representation of the average cost of overflows that the collector would pay at the 75th, 90th, 95th and 99th percentile of the simulated scenarios over all 63 instances, for the routing only solution and for the complete solution with an ECC of 100 CHF and an RFCM equal to one. The differences are consequential. The cost due to the routing-only objective is from 20 times higher at the 75th percentile to 8 times higher at the 99th percentile, which is a clear indication of the underestimation of risk in the face of stochastic demand. Even at the 99th percentile, the complete objective would result, on average, in a total cost

Table 3.8: Probabilistic Policies: Container Overflows and Route Failures for Real Data Instances

Objective	ECC	RFCM	Avg Num Overflows				Avg Num Route Failures			
			75th Perc.	90th Perc.	95th Perc.	99th Perc.	75th Perc.	90th Perc.	95th Perc.	99th Perc.
Complete	100.00	1.00	0.83	1.60	2.15	3.26	0.05	0.05	0.05	0.07
Complete	100.00	0.50	0.81	1.58	2.14	3.27	0.05	0.06	0.07	0.10
Complete	100.00	0.25	0.81	1.59	2.15	3.26	0.05	0.07	0.07	0.11
Complete	100.00	0.00	0.83	1.57	2.16	3.28	0.10	0.11	0.12	0.16
Complete	50.00	1.00	1.04	1.87	2.48	3.72	0.05	0.05	0.05	0.05
Complete	50.00	0.50	1.04	1.86	2.48	3.73	0.05	0.07	0.07	0.07
Complete	50.00	0.25	1.06	1.88	2.50	3.72	0.06	0.09	0.09	0.10
Complete	50.00	0.00	1.06	1.87	2.48	3.72	0.09	0.11	0.11	0.13
Complete	25.00	1.00	1.26	2.12	2.73	4.08	0.06	0.06	0.06	0.06
Complete	25.00	0.50	1.25	2.10	2.73	4.07	0.05	0.07	0.07	0.07
Complete	25.00	0.25	1.25	2.11	2.74	4.09	0.05	0.08	0.08	0.09
Complete	25.00	0.00	1.25	2.11	2.77	4.09	0.09	0.10	0.11	0.11
Routing-only	0.00	0.00	16.93	20.45	22.55	26.71	0.04	0.05	0.05	0.05

of less than 1000 CHF, compared to more than 3000 CHF for the routing-only objective.

Figure 3.4 is a box-plot of the average number of simulated container overflows for the 63 instances at the 75th, 90th, 95th and 99th percentile. For the complete objective with an ECC of 100 CHF and an RFCM equal to one, the ratio of container overflows to the number of containers in the instance goes from 2% at the 75th to 9% at the 99th percentile. For the routing-only objective, these values are 41% and 66%, respectively.

To study the main drivers of the number of container overflows, we perform a series of linear regressions. Table 3.9 consists of two parts. In part (a) the explanatory factor is the forecasting error ζ , while in part (b) it is the number of containers in the instance. In both parts, the first column identifies the type of objective considered, and the rest of the columns correspond to the dependent variable, i.e. the average number of container overflows over the 63 instances at the 75th, 90th, 95th and 99th percentile. For each of them, we report the coefficient of the explanatory factor followed by a significance code, and the coefficient of determination R^2 . We observe that all coefficients are positive as expected. The regressions on the forecasting error suggest that it explains approximately half of the variability in the container overflows for the routing-only objective and about 40% in the case of the complete objective with an ECC of 100 CHF and an RFCM equal to one. This result is intuitive as higher forecasting errors lead to larger demand perturbations in the simulation experiments and, as

Figure 3.3: Average Cost of Overflows at Different Percentiles of the Simulated Scenarios




Figure 3.4: Overflows for All Instances at Different Percentiles of the Simulated Scenarios




Table 3.9: Driving Factors for the Occurrence of Container Overflows

Objective	(a) Regressions on Forecasting Error ζ							
	75th Percentile		90th Percentile		95th Percentile		99th Percentile	
	Coefficient	R^2	Coefficient	R^2	Coefficient	R^2	Coefficient	R^2
Complete with ECC=100, RFCM=1	0.02***	0.41	0.02***	0.35	0.03***	0.45	0.03***	0.46
Routing-only	0.16***	0.52	0.18***	0.53	0.19***	0.52	0.21***	0.51

Objective	(b) Regressions on Number of Containers in the Instance							
	75th Percentile		90th Percentile		95th Percentile		99th Percentile	
	Coefficient	R^2	Coefficient	R^2	Coefficient	R^2	Coefficient	R^2
Complete with ECC=100, RFCM=1	0.02	0.03	0.01	0.01	0.02	0.02	0.02	0.01
Routing-only	0.34***	0.24	0.37***	0.24	0.40***	0.25	0.46***	0.26

Note. Significance codes: *** 99%.

a consequence, to a higher rate of overflows. Nevertheless, the coefficient values in the case of the complete objective with an ECC of 100 CHF and an RFCM equal to one are very low, suggesting that the forecasting error has only a very slight effect on the number of overflows. The results of the regressions on the number of containers in the instance exhibit an even more pronounced difference. While it can explain approximately 25% of the variability in the container overflows for the routing-only objective, the number of containers in the instance seems not to have an effect on the overflows in the case of the complete objective with an ECC of 100 CHF and an RFCM equal to one. We observe no significant coefficients in the latter case. This is a desirable result as it would suggest that the number of overflows does not scale up with the instance size. It also has a managerial implication, giving a reliable estimate of extreme events over a wide range of situations.

Appendix A provides additional analysis of complementary results. In particular, Appendix A.1 explores the differences in collection strategies between the complete objective solutions and the routing-only solutions, showing that the complete objective anticipates overflows during the weekend when vehicles are not available and collects containers preventively although they are relatively empty. Appendix A.2 studies the effect of using smaller capacity vehicles on the route failure cost, showing that the ERFC component remains a relatively small fraction of the total cost.

Alternative Policies.

To further study the theoretical justification and practical relevance of the probabilistic approach, we compare it to an intuitive routing-only approach, in which during the solution of the problem we use artificially low capacities for the containers and the trucks. This policy is an attempt to control the number of container overflows and route failures and it also leads, undoubtedly, to higher routing costs due to the necessity of more frequent visits. After each instance is solved, we perform the same simulation-based validation of the solution as before. However, during the simulation we count the number of container overflows and

Chapter 3. The Waste Collection IRP with Stochastic Demands

route failures with respect to the original container and truck capacities. Thus, we have a fair comparison between the probabilistic approach and the alternative policies of artificially low capacities.

Tables 3.10, 3.11 and 3.12 are structured in the same way as Tables 3.6, 3.7 and 3.8. Here, the objective is always routing-only and what varies are the Container Effective Capacity (CEC) and the Truck Effective Capacity (TEC) as fractions of their original capacities. In Table 3.10, we note the strong negative correlation between the container effective capacity and the average number of tours, container collections and dump visits in the solutions. We also notice that the relative increase in the number of container collections is much higher than the reduction of the container effective capacity. This is an artifact of the finite planning horizon as many containers may be collected two or three times rather than once or twice due to their smaller effective capacities. This effect will most likely diminish over the long run. We notice that the solution time grows with the number of container collections, and so do the solution gaps. Yet, the increase of the solution time is smaller than the increase of the number of container collections. Moreover, even the highest gaps for a container effective capacity of 60% remain in the order of 1.5% and below. One explanation for the increase of solution time and the gaps could be that the problem becomes tighter and hence the solution space more challenging. In fact, two of the instances for a CEC of 60% are infeasible.

Table 3.11 shows the gradual growth of the routing cost as we reduce the effective capacities. Since the objective is always routing-only, the overflow and route failure components do not apply. The last three columns reveal an interesting result. Lowering the CEC from 100%, to 90%, to 75% results in solutions collecting more volume, but also more volume per unit routing cost. However, further lowering the container effective capacity to 60% results in a disproportionately higher routing cost. As a result, despite collecting more volume, the solutions with a CEC of 60% are less efficient in terms of collected volume per unit routing cost compared to the solutions with a CEC of 75%. Table 3.12 describes the average results of the 10,000 simulation scenarios that were performed on each instance with the original

Table 3.10: Alternative Policies: Basic Results for Cost Analysis on Real Data Instances

Objective	CEC	TEC	Runtime(s.)	Avg Num Tours	Avg Num Collections	Avg Num Dump Visits	Best Cost (CHF)	Avg Cost (CHF)	Worst Cost (CHF)	Gap Avg- Best(%)	Gap Worst- Best(%)
Routing-only	1.00	1.00	682.31	1.83	16.64	1.87	421.95	422.51	423.16	0.13	0.29
Routing-only	1.00	0.90	685.38	1.83	16.65	1.87	422.22	422.80	423.47	0.14	0.30
Routing-only	1.00	0.75	672.96	1.83	16.65	1.95	423.38	424.02	424.92	0.15	0.36
Routing-only	1.00	0.60	757.33	1.83	16.66	2.04	425.31	426.06	426.93	0.18	0.38
Routing-only	0.90	1.00	742.70	2.00	22.63	2.02	486.29	486.83	487.59	0.11	0.27
Routing-only	0.90	0.90	746.77	2.00	22.62	2.06	486.82	487.39	488.09	0.12	0.26
Routing-only	0.90	0.75	738.18	2.00	22.62	2.15	488.46	489.16	489.95	0.14	0.31
Routing-only	0.90	0.60	725.43	2.00	22.63	2.37	492.74	493.71	494.69	0.20	0.39
Routing-only	0.75	1.00	873.54	2.00	33.52	2.43	541.87	542.92	544.53	0.19	0.49
Routing-only	0.75	0.90	863.36	2.00	33.52	2.60	544.60	545.78	547.25	0.22	0.49
Routing-only	0.75	0.75	869.94	2.00	33.50	2.86	549.13	550.15	551.46	0.19	0.42
Routing-only	0.75	0.60	862.67	2.00	33.54	3.12	555.35	557.37	559.75	0.36	0.79
Routing-only	0.60	1.00	1037.72	2.97	44.59	3.78	780.40	783.05	788.46	0.34	1.03
Routing-only	0.60	0.90	1241.91	2.97	44.65	3.88	782.50	785.42	792.24	0.37	1.25
Routing-only	0.60	0.75	1060.95	2.97	44.67	4.10	788.74	792.06	798.07	0.42	1.18
Routing-only	0.60	0.60	1023.95	2.97	44.79	4.58	799.71	804.37	811.70	0.58	1.50

3.4. Numerical Experiments

container and truck effective capacities. It is immediately clear that considering artificially low capacities during the solution has a marked effect in reducing overflows and route failures. To be precise, the average number of overflows drops by roughly a third when the container effective capacity is reduced to 90% and by roughly two thirds when it is reduced to 75%. On the other hand, reducing the truck effective capacity to 90% can effectively eliminate the occurrence of route failures.

Figures 3.5 and 3.6 present a side-by-side comparison of the probabilistic and the alternative policies of using artificially low container and truck capacities. In both figures, the first 12 bars represent the probabilistic model with complete objective function for various Emergency Collection Costs (ECC) and Route Failure Cost Multipliers (RFCM). The last 16


Table 3.11: Alternative Policies: Key Performance Indicators for Cost Analysis on Real Data Instances

Objective	CEC	TEC	Avg Routing Cost (CHF)	Avg Overflow Cost (CHF)	Avg Rte Failure Cost (CHF)	Avg Collected Volume (L)	Liters Per Unit Cost	Liters Per Unit Routing Cost
Routing-only	1.00	1.00	422.51	0.00	0.00	24,992.02	59.15	59.15
Routing-only	1.00	0.90	422.80	0.00	0.00	24,963.64	59.04	59.04
Routing-only	1.00	0.75	424.02	0.00	0.00	24,986.17	58.93	58.93
Routing-only	1.00	0.60	426.06	0.00	0.00	24,909.59	58.46	58.46
Routing-only	0.90	1.00	486.83	0.00	0.00	31,553.37	64.81	64.81
Routing-only	0.90	0.90	487.39	0.00	0.00	31,577.74	64.79	64.79
Routing-only	0.90	0.75	489.16	0.00	0.00	31,747.19	64.90	64.90
Routing-only	0.90	0.60	493.71	0.00	0.00	31,846.97	64.51	64.51
Routing-only	0.75	1.00	542.92	0.00	0.00	44,149.46	81.32	81.32
Routing-only	0.75	0.90	545.78	0.00	0.00	44,108.02	80.82	80.82
Routing-only	0.75	0.75	550.15	0.00	0.00	43,985.69	79.95	79.95
Routing-only	0.75	0.60	557.37	0.00	0.00	44,219.61	79.34	79.34
Routing-only	0.60	1.00	783.05	0.00	0.00	54,332.98	69.39	69.39
Routing-only	0.60	0.90	785.42	0.00	0.00	54,360.53	69.21	69.21
Routing-only	0.60	0.75	792.06	0.00	0.00	54,479.13	68.78	68.78
Routing-only	0.60	0.60	804.37	0.00	0.00	54,564.10	67.83	67.83

Table 3.12: Alternative Policies: Container Overflows and Route Failures for Real Data Instances

Objective	CEC	TEC	Avg Num Overflows				Avg Num Route Failures			
			75th Perc.	90th Perc.	95th Perc.	99th Perc.	75th Perc.	90th Perc.	95th Perc.	99th Perc.
Routing-only	1.00	1.00	16.96	20.47	22.58	26.71	0.03	0.05	0.05	0.05
Routing-only	1.00	0.90	16.93	20.42	22.54	26.68	0.00	0.00	0.00	0.00
Routing-only	1.00	0.75	16.90	20.42	22.55	26.70	0.00	0.00	0.00	0.00
Routing-only	1.00	0.60	16.85	20.37	22.50	26.63	0.00	0.00	0.00	0.00
Routing-only	0.90	1.00	10.29	13.07	14.78	18.23	0.02	0.02	0.02	0.02
Routing-only	0.90	0.90	10.25	13.04	14.74	18.15	0.00	0.00	0.00	0.00
Routing-only	0.90	0.75	10.27	13.03	14.77	18.15	0.00	0.00	0.00	0.00
Routing-only	0.90	0.60	10.28	13.02	14.77	18.21	0.00	0.00	0.00	0.00
Routing-only	0.75	1.00	4.23	6.07	7.25	9.65	0.06	0.06	0.06	0.06
Routing-only	0.75	0.90	4.25	6.06	7.27	9.66	0.00	0.00	0.00	0.00
Routing-only	0.75	0.75	4.25	6.07	7.29	9.68	0.00	0.00	0.00	0.00
Routing-only	0.75	0.60	4.24	6.03	7.25	9.67	0.00	0.00	0.00	0.00
Routing-only	0.60	1.00	2.17	3.52	4.45	6.34	0.01	0.01	0.01	0.01
Routing-only	0.60	0.90	2.18	3.52	4.48	6.32	0.00	0.00	0.00	0.00
Routing-only	0.60	0.75	2.15	3.54	4.46	6.29	0.00	0.00	0.00	0.00
Routing-only	0.60	0.60	2.17	3.53	4.47	6.31	0.00	0.00	0.00	0.00


Figure 3.5: Comparison of Routing Cost for Probabilistic and Alternative Policies



bars represent the alternative policies of using artificially low capacity for various Container Effective Capacities (CEC) and Truck Effective Capacities (TEC). We should point out that the baseline routing-only policy with container and truck effective capacity of 100% has the lowest routing cost. Figure 3.5 reveals that the routing cost of the probabilistic policies considered ranges from approximately 550 to 580 CHF depending mostly on the value of the emergency collection cost. This latter range is relatively limited compared to the range of routing costs for the alternative policies, which goes from 420 to 800 CHF, with pronounced jumps linked to the variation of the container effective capacity. We observe a disproportionate cost increase linked to lowering the container effective capacity from 75% to 60%. This effect is due to the fact that many more containers need to be collected now. There are on average three vs. two tours per solution, compared to the case of a CEC of 75% or 90%. Moreover, tours are on average also longer and as a result less compact.

We contrast the above observation with the average number of overflows and route failures after the simulation-based validation of both types of policies. These are presented in Figure 3.6, in parts (a) and (b), respectively. What part (a) of the figure reveals is that all considered probabilistic policies are able to contain the number of overflows to very low values. There is still a slight increase in the number of overflows (with an associated slight decrease in the routing cost) when the emergency collection cost is reduced from 100 to 50 and then to 25 CHF. Nevertheless, the average number of overflows across all instances is approximately four, even at the 99th percentile. In comparison, the average number of overflows for the alternative policies is markedly higher. While reducing the container effective capacity leads to a considerable drop in the number of overflows, the alternative policies cannot beat the probabilistic ones. A case in point are the complete objective solutions for an ECC of 25 CHF and the routing-only solutions for a CEC of 75%. While they incur a similar routing cost as shown in Figure 3.5, Figure 3.6 reveals that the occurrence of overflows for the routing-only solutions is more than twice as high. Reducing further the CEC leads to a mild decrease in the occurrence of overflows accompanied by a significant increase in the routing cost. We must stress here that since we compare the performance of

Figure 3.6: Comparison of Container Overflows and Route Failures for Probabilistic and Alternative Policies



two policy types in terms of number of overflows and route failures at different percentiles, we must isolate these components from the solution cost of the probabilistic model. The above findings clearly indicate the superior performance of the probabilistic policies in the face of stochastic demand. Whereas the alternative policies can only control overflows in the expected sense, the probabilistic model attributes a cost to this uncertainty over the whole planning horizon. Thus, it uses foresight in a much more intelligent way.

Lastly, part (b) of Figure 3.6 shows how both types of policies perform in terms of the average number of route failures over all instances. Here, the alternative policies appear to be more successful. As already noted before, reducing the truck effective capacity to 90% is sufficient to eliminate the occurrence of route failures. As far as the probabilistic policies are concerned, we identify an interesting pattern. The number of route failures is positively correlated with the emergency collection cost and negatively correlated with the route failure cost multiplier. The latter is an intuitive result. The former relationship, however, is slightly more intricate. What is at play here is a trade-off between container overflows and route failures. A higher emergency collection cost incentivizes more frequent visits. Trucks thus

collect more containers in each tour and, by consequence, in each depot-to-dump or dump-to-dump visit. Since trucks are fuller on average, the solution is subject to a higher risk of route failures. The probabilistic policies collect on average more containers than the alternative policies and this could be a valid explanation of the latter's better performance when it comes to limiting the number of route failures. However, as reported in Table 3.7, the contribution of the expected route failure cost to the total solution cost for our particular instances is in any case immaterial.

Rolling Horizon Approach.

In practice, the SIRP that we consider will be solved on a daily rolling horizon basis using the latest available container level information. In this approach, the problem is solved for a planning horizon \mathcal{T} , the tours that are scheduled on day $t = 0$ are executed, the horizon is rolled over by a day, the problem is re-solved for the new initial container levels and updated forecasts, and so on. Thus, true demands are gradually revealed each day, but the demands over the planning horizon are still stochastic. This type of problem is known as the Dynamic and Stochastic Inventory Routing Problem (DSIRP). The solution of the DSIRP requires the solution of an SIRP at each rollover. The cost of the DSIRP is composed of the total routing and overflow cost on day $t = 0$ resulting from the solution of the SIRP at each rollover. We note that the route failure cost does not apply on day $t = 0$. We also note that overflows on day $t = 0$ are deterministic, since the container levels are fully known, and thus for each overflow on day $t = 0$ the full overflow cost χ is paid.

In the solution of the DSIRP, true demands are gradually revealed in the solution process, which reduces uncertainty. Thus, we hypothesize that its solution cost should be bounded above by the solution cost of a static SIRP for the same planning horizon. Assume that we solve the SIRP for a planning horizon $\mathcal{T} = \{0, \dots, u\}$. In order to compare its cost to that of the DSIRP, we should roll over for a number of times equal to the length of the planning horizon \mathcal{T} , i.e. the last rollover should be on day u . Moreover, for rollover t the initial container levels are updated by true demands and also dependent on the solution of rollover $t - 1$. Updated forecasts should ideally be used if available. We also hypothesize that the solution of the DSIRP should be bounded below by the solution of a static IRP using true demands for the same planning horizon \mathcal{T} . Using true demands rids the problem of any uncertainty. The solution of the IRP results in an intelligent assignment of tours to days. Thus, the number of executed tours over the planning horizon will be minimized and tours may not be executed on each day. This is not necessarily the case for the solution of the DSIRP, which has no memory of the past rollovers and may assign tours on day $t = 0$ for each rollover.

To test our hypotheses, we generate a second set of real data instances. It comprises 41 instances, each covering two weeks of white glass collections in the canton of Geneva, Switzerland in 2014, 2015, or 2016. On average, there are 69 containers per instance, and the maximum is 86. Otherwise, the instances fit the same description as the set of 63 instances

Table 3.13: Analysis of Rolling Horizon DSIRP Bounds

Instance	Static IRP with True Demand	Rolling DSIRP with Forecast Demand	Static SIRP with Forecast Demand	Instance	Static IRP with True Demand	Rolling DSIRP with Forecast Demand	Static SIRP with Forecast Demand
Inst_1	276.44	585.69	658.39	Inst_22	429.20	526.06	607.22
Inst_2	448.67	937.47	849.43	Inst_23	241.44	568.15	681.54
Inst_3	307.88	626.01	816.05	Inst_24	547.92	769.08	747.64
Inst_4	266.15	577.82	701.61	Inst_25	446.31	583.87	689.37
Inst_5	450.14	663.50	790.44	Inst_26	442.02	575.57	656.27
Inst_6	300.73	624.62	708.79	Inst_27	441.36	595.47	705.01
Inst_7	268.65	580.83	649.67	Inst_28	465.74	628.59	733.80
Inst_8	427.17	608.31	680.36	Inst_29	436.25	579.74	701.33
Inst_9	442.34	609.44	656.44	Inst_30	414.41	701.87	692.33
Inst_10	448.70	578.34	647.05	Inst_31	442.87	530.14	668.17
Inst_11	467.35	614.28	669.33	Inst_32	255.32	617.04	695.62
Inst_12	449.20	681.10	625.59	Inst_33	460.04	641.00	773.33
Inst_13	254.66	558.57	629.36	Inst_34	505.55	674.98	710.84
Inst_14	276.60	613.72	685.64	Inst_35	481.85	746.10	786.94
Inst_15	429.26	562.12	788.75	Inst_36	454.60	658.51	741.02
Inst_16	529.60	626.97	702.61	Inst_37	465.33	651.41	749.50
Inst_17	423.07	589.66	663.90	Inst_38	519.56	709.76	809.91
Inst_18	457.65	596.14	681.29	Inst_39	243.94	623.29	697.93
Inst_19	448.66	524.41	596.81	Inst_40	450.94	620.09	756.48
Inst_20	418.12	569.73	653.22	Inst_41	403.01	576.45	688.68
Inst_21	276.32	570.41	622.47				

Note. The four instances for which the hypothesized bounds do not hold are shown in bold.

described in the beginning of the section. We solve the static IRP with true demands and static SIRP with forecast demands for the first week, and the DSIRP with a one week planning horizon and rollovers for the first week. Table 3.13 presents the results we obtain. Since we are interested in verifying the empirical existence of the hypothesized bounds, we report the best cost out of 10 runs for each instance. We observe that the hypothesized bounds are obtained in all but four cases, which are shown in bold. The relative differences are also very interesting to look at. The solutions of the DSIRP are on average 61% more expensive than those of the static IRP with true demands. This result is inevitably related to the level of uncertainty as represented by the forecasting error ζ . In other words, if more accurate forecasting methodologies are available, this gap may be brought down. On the other hand, the static SIRP approach is on average 14% more expensive than the rollover approach for the DSIRP. This clearly shows the benefit of applying the latter in practical situations.

3.5 Summary

In this chapter, we motivate and formulate the waste collection IRP, an extension of the waste collection VRP from Chapter 2. Demand is stochastic and non-stationary, and can be forecast by any model that provides the expected demands over the planning horizon and a measure of uncertainty represented by the standard deviation of the error terms, the latter assumed to be iid normal. This results in a rich stochastic inventory routing problem, whose objective captures demand uncertainty with the goal of minimizing the expected cost, including the expected cost of recourse actions, subject to a range of practical and policy-related constraints. To solve the problem, we develop an ALNS algorithm which produces excellent results on IRP benchmarks sets from the literature, as well as very good

Chapter 3. The Waste Collection IRP with Stochastic Demands

results on rich vehicle routing instances. The application of the methodology to instances derived from real data confirms the relevance of the probability information in the objective function.

The computational experiments demonstrate that including probabilistic information in the objective function leads to a relatively modest increase in the routing cost, while avoiding major expenditures that otherwise occur even at moderate percentiles of the simulated demand scenarios. Based on our policy, we can control the rate of occurrence of undesirable events, like overflows and route failures, by scaling the probability-related costs considered in the objective function. The probabilistic approach significantly outperforms alternative policies of using artificially low capacities for the containers and the trucks in its ability to control the occurrence of container overflows for the same routing cost. We also analyze the solution properties of a rolling horizon approach and verify the hypothesized lower and upper bounds on the solution cost.

Chapter 4 extends the ideas presented here to a unified framework for rich routing problems with stochastic demands. It relaxes fully or partially the assumption of iid normal error terms, considers a more general inventory policy, discusses tractability related topics, and illustrates applications to a variety of rich routing problems borrowed from the literature and inspired from practice.

4 A Unified Framework for Rich Routing Problems with Stochastic Demands

This chapter is an extension of the article:

Markov, I., Bierlaire, M., Cordeau, J.-F., Maknoon, Y., and Varone, S. (2017). A general framework for routing problems with stochastic demands. *Proceedings of the 17th Swiss Transport Research Conference (STRC)*, May, 17-19, 2017.

The work therein has been performed by the author in collaboration with Prof. Michel Bierlaire, Prof. Jean-François Cordeau, Prof. Yousef Maknoon and Prof. Sacha Varone.

This chapter generalizes the waste collection IRP of Chapter 3 in a unified framework for modeling and solving rich routing problems with stochastic demands, including among others the VRP and the IRP. We relax fully or partially the assumption of iid normal error terms and discuss the ensuing tractability issues and how we tackle them.

The two classical IRP inventory policies are the Order-Up-to (OU) level policy and the Maximum Level (ML) policy (Bertazzi et al., 2002; Archetti et al., 2011). Under the former delivery is up to the capacity ω_i , while under the latter the delivery quantity is part of the decisions. We consider a discretized ML policy, which is more general than the OU policy used in the waste collection IRP. The reason for this discretization is again tractability, as the presence of discrete levels allows us to pre-process much of the probability information before starting the solution process.

We extend and refine the notation of Chapter 3, and propose a general optimization model which can be adapted to various distribution and collection VRP and IRP problems from the literature and practice, from waste collection to maritime inventory routing. To assess the suitability of the framework, we perform numerical experiments on the waste collection IRP instances of Chapter 3 and on a new set of instances for the facility maintenance problem, in which the probability of facility breakdowns accumulates with time in a way similar to inventory.

Chapter 4. A Unified Framework for Rich Routing Problems with Stochastic Demands

The chapter is organized as follows. Section 4.1 offers a brief review of the relevant literature on rich routing problems from several application areas, including health care, waste collection, and maritime, with a focus on demand stochasticity. Section 4.2 introduces the main concepts and modeling elements used by the unified framework. These are further discussed and elaborated in Section 4.3, which details the treatment of demand stochasticity, and Section 4.4, which develops the optimization model. In turn, Section 4.5 provides examples of adapting the framework to various specific problem classes. Section 4.6 presents the numerical experiments and, finally, Section 4.7 ends with a summary of the main findings and contributions.

4.1 Related Literature

This section offers a literature review of routing problems with stochastic demands, starting from rich vehicle and inventory routing problems in general and then exploring several specific and pertinent application areas. The analysis comments on the variety of approaches used in integrating stochastic demand in the modeling and solution process, thus highlighting the need for a unified approach.

4.1.1 Rich Vehicle and Inventory Routing Problems

Rich vehicle routing problems are multi-constrained routing problems that extend the classical capacitated VRP (Dantzig and Ramser, 1959) by including a variety of features relevant to real-world problems, such as heterogeneous fleets, time windows, driver constraints, open tours, multiple depots or intermediate replenishment facilities, dynamism, stochastic information, etc. The recent work of Lahyani et al. (2015) develops a taxonomy and a definition of rich VRPs. Surveys on various aspects concerning heterogeneous fleets, intermediate replenishment facilities, time windows, open tours and multiple depots are available in Markov et al. (2014) and Chapter 2. Rich routing problems often include an uncertainty component. In dynamic problems, parameters are partly unknown and gradually revealed with time. In dynamic and stochastic problems, we have access to probability information of the unknown parameters. Ritzinger et al. (2016) summarize the recent literature on dynamic and stochastic VRPs and offer a classification scheme based on the available stochastic information. Gendreau et al. (2016) center their survey on the state of the art of the a priori and the re-optimization paradigms for stochastic VRPs, the two being the predominantly used paradigms by researchers.

Although multi-constrained IRPs with real-world features have recently begun to appear in the literature, the term rich IRP has not established itself as in the case of the VRP. Zhalechian et al. (2016) and Soysal (2016), for example, discuss closed-loop IRP systems with stochastic demands. Both include environmental considerations in the objective function. Zhalechian et al. (2016) also include social considerations, present a fuzzy approach, and develop a

hybrid meta-heuristic and a lower bounding procedure, which are applied on a small case study. Soysal (2016) use CPLEX to solve a small case study and, based on a simulation experiment, confirm the benefit of including uncertainty in the model. Rahimi et al. (2017) describe a rich IRP with environmental considerations and stochastic parameters, including stochastic demand, and propose a fuzzy approach. Their solution methodology relies on a meta-heuristic from the literature. However, the focus of their numerical experiments is not on the effect of uncertainty. Furthermore, none of these studies models explicitly recourse actions in the events of stock-outs and route failures, which occur as a consequence of demand uncertainty. Chapter 1 discussed the advantages and disadvantages of various modeling approaches, while Chapter 3 provided a review of the literature on road-based stochastic IRPs with a finite horizon dimension. Sections 4.1.2, 4.1.3 and 4.1.4 below extend the survey to several additional application areas of routing problems with stochastic demands that can be modeled using the unified framework. Then, Section 4.1.5 positions our approach.

4.1.2 Health Care Routing Problems

Stochastic demand appears in health care routing problems involving the pick-up and delivery of drugs, biological samples, and medical equipment. Hemmelmayr et al. (2010) solve a stochastic blood distribution problem, which considers shortfalls and spoilage. To balance delivery and spoilage costs, they limit the probability of spoilage to 5% by sampling product usage during the spoilage period and taking the 5% quantile as the maximum inventory level at the hospital. Hemmelmayr et al. (2010) develop a two-stage stochastic program with recourse, assuming knowledge of the inventory in the beginning of each day of the planning horizon. The authors extend an exact approach and a VNS meta-heuristic from the literature, in both cases using external sampling to convert the two-stage stochastic optimization problem into a deterministic one. Through a simulation experiment, they show that a simple recourse policy is sufficient to provide a reliable and cost-efficient blood supply. Niakan and Rahimi (2015) and Shi et al. (2017) study the problem of delivering drugs with uncertain demands to patient homes. Both authors apply fuzzy programming approaches to the problem and report the added value of incorporating uncertainty into the model. The broader literature on health care routing problems identifies workload balancing and the continuity of service, or continuity of care in this specific context, as two of the most important concerns in this field (see e.g. Lanzarone and Matta, 2009, 2012; Lanzarone et al., 2012; Errarhout et al., 2014, 2016).

4.1.3 Waste Collection Routing Problems

Chapter 3 describes a stochastic IRP for the collection of recyclable waste with the integration of demand forecasting. Demand stochasticity leads to the occurrence of container overflows and route failures. The proposed stochastic model significantly outperforms alternative

deterministic policies in its ability to limit the occurrence of container overflows for the same routing cost. Still in the area of waste collection, Johansson (2006) and Mes (2012) use simulation to confirm the benefits of migrating from static to dynamic collection policies in Malmö, Sweden and a study area in the Netherlands, respectively, where containers are equipped with level and motion sensors, respectively. Mes (2012) finds a positive added value of investing in level sensors compared to simple motion sensors that detect when a container is emptied. Mes et al. (2014) apply optimal learning techniques to tune the parameters related to inventory control (deciding which containers to select) assuming accurate container level information. Nolz et al. (2011) develop a tabu search algorithm for a stochastic IRP for the collection of infectious waste from pharmacies. Nolz et al. (2014b) propose a scenario sampling method and an ALNS algorithm for the same problem. Nolz et al. (2014a) extend this to a bi-objective problem, trading off satisfaction of pharmacies, local authorities and the minimization of public health risks against routing costs. They propose three meta-heuristic approaches for this problem. Bitsch (2012) develops a VNS for an IRP applied to the collection of recyclable waste in a Danish region. Waste level is stochastic and containers should be emptied so that the probability of overflow is six standard deviations away.

4.1.4 Maritime Routing Problems

Papageorgiou et al. (2014) identify three features that distinguish maritime from road-based IRPs, specifically: 1) the absence of a central depot, which entails multi-period open tours, 2) the long travel times and port operations, which prolong the planning horizon, and 3) the shorter succession of port visits, in comparison to the typically dozens of customer visits in road-based IRPs. Cheng and Duran (2004) solve a crude oil transportation problem with inventory management, integrating discrete event simulation and stochastic optimal control. The optimal control problem is formulated as a Markov decision process that incorporates travel time and demand uncertainty. Yu (2009) discusses a problem with multiple supply and demand ports, where the only stochastic element is the demand. It is formulated as a stochastic program and branch-and-price is used to solve medium-sized instances. Arslan and Papageorgiou (2015) study a maritime fleet renewal and deployment problem under demand and charter cost uncertainty, which determines the fleet size, mix, and deployment strategy to satisfy stochastic demands over the planning horizon. They solve the problem in a rolling horizon fashion using a stochastic programming look-ahead model, and explore the impact of different scenario trees with different recourse functions. Zheng and Chen (2016) propose a real option model to solve a fleet replacement model under demand and fuel price uncertainty. Monte Carlo simulation is used to find replacement probabilities in future years and the net present value of cost savings. The distribution of Liquefied Natural Gas (LNG) is a particularly important application area. Moraes and Faria (2016) study an LNG planning problem for an oil and gas company. They develop a two-stage stochastic linear model to address uncertainties related to the LNG demand and spot prices. Halvorsen-Weare et al. (2013) consider an LNG routing and scheduling problem with time windows, berth capacity

and inventory level constraints. They propose and test various robustness strategies with respect to travel times and daily LNG production rates.

4.1.5 Discussion

The reviewed literature employs a variety of approaches for capturing demand uncertainty, with different simplifying assumptions and modeling techniques, and with or without explicit recourse policies and penalties for the occurrence of undesirable events. As discussed in Chapter 1, all these approaches come with their advantages and limitations. We identified our goal of modeling the cost of demand uncertainty and bridging the gap between theory and practice, preserving at the same time solution tractability for realistic-size problems.

In this chapter, we develop a unified modeling and solution framework for rich routing problems with stochastic demands. Using a set of key concepts and modeling elements, it provides a common language for describing and modeling such problems and imposes very few distributional assumptions on the demand. The approach distinguishes itself through several unifying features, namely 1) the applicability to various types of rich routing problems, including among others rich VRPs and IRPs, 2) the integration of real demand forecasting without a stationarity assumption, 3) the inclusion of the cost of undesirable events, such as stock-outs/overflows and route failures, in the objective function, 4) the explicit modeling of the cost of recourse actions in response to the above events, 5) the tractability of the resulting framework due to the ability to pre-compute or at least partially pre-process most of the stochastic information for a general inventory policy, 6) and the intuitive evaluation of the produced solution through simulation. Simulation is used both to measure the frequency of occurrence of undesirable events in the final solution and to evaluate how closely it models the real cost given the imposed assumptions. Integrating all cost information in the objective function, the resulting solution reflects the actual cost of demand uncertainty.

4.2 Key Concepts and Modeling Elements

This section introduces the key concepts, the modeling elements, such as sets and parameters, and the relationships among them, which are used in the development and formulation of the unified framework in Sections 4.3 and 4.4 below. For the sake of generality of presentation, consider a problem in a distribution context. We comment on the changes that apply to a collection context when needed. Building on the notation established in Chapter 3, we consider a planning horizon $\mathcal{T} = \{0, \dots, u\}$ of discrete time periods, such as days or another appropriate level of discretization. Deliveries are performed by a heterogeneous fixed fleet \mathcal{K} , with each vehicle $k \in \mathcal{K}$ defined by a per-period deployment cost φ_k , a unit-distance running cost β_k , a unit-time running cost θ_k , and a capacity Ω_k . The fleet reduces to a homogeneous one if the values of these parameter are identical for all vehicles.

Chapter 4. A Unified Framework for Rich Routing Problems with Stochastic Demands


For each vehicle $k \in \mathcal{K}$ and period $t \in \mathcal{T}$, we are given a directed graph $\mathcal{G}_{kt}(\mathcal{N}_{kt}, \mathcal{A}_{kt})$. The set \mathcal{O} includes all origin and destination depots, where $\mathcal{O}'_{kt} \subseteq \mathcal{O}$ is the set of origin depots for vehicle k in period t and $\mathcal{O}''_{kt} \subseteq \mathcal{O}$ is the set of destination depots for vehicle k in period t . In addition, \mathcal{P} is the set of demand points, \mathcal{D} is the set of supply points, $\mathcal{N}_{kt} = \mathcal{O}'_{kt} \cup \mathcal{O}''_{kt} \cup \mathcal{P} \cup \mathcal{D}$ is the set of all points potentially reachable by vehicle k in period t , and $\mathcal{A}_{kt} = \{(i, j) : \forall i, j \in \mathcal{N}_{kt}, i \neq j\}$ is the set of arcs connecting the latter. The set \mathcal{D} contains a sufficient number of replications of each supply point to allow multiple visits by the same vehicle in the same period. The distance matrix is asymmetric, with π_{ij} the length of arc $(i, j) \in \mathcal{A}_{kt}$, for any vehicle k and period t . Vehicle k can have a specific travel time matrix for each period t , where τ_{ijk_t} is the travel time of vehicle k on arc $(i, j) \in \mathcal{A}_{kt}$ in period t . Point $i \in \mathcal{O} \cup \mathcal{P} \cup \mathcal{D}$ presents a time window $[\lambda_i, \mu_i]$, where λ_i and μ_i stand for the earliest and latest possible start-of-service time at that point. Start of service after μ_i is not allowed and if the vehicle arrives before λ_i , it has to wait. Service duration at point i is denoted by δ_i , with service durations in the set \mathcal{O} being zero.

With each demand point $i \in \mathcal{P}$ is associated an inventory capacity of ω_i , a visit cost of ξ_i , and an inventory holding cost of η_i . The parameter ν_i specifies the minimum number of times that demand point i must be visited over the planning horizon. There is the option of imposing periodicity on the visits as well. The set \mathcal{C}_i contains the visit period combinations for demand point i , and the binary constant α_{rt} denotes whether period t belongs to visit period combination $r \in \mathcal{C}_i$ for demand point i . The binary flags α_{ikt} denote whether point $i \in \mathcal{P} \cup \mathcal{D}$ is accessible by vehicle k in period t . They can also be used to express continuity of service, restricting the vehicle(s) that can visit demand point i .

In period t , demand point i exhibits non-stationary stochastic demand ρ_{it} . It is important to highlight that stochasticity refers to normal operations, and not to hazard or deep uncertainty (Gendreau et al., 2016). Demand stochasticity implies a probability of stock-out, one of two possible states for each demand point, which happens when its inventory becomes negative. Let $\sigma_{it} = 1$ denote that demand point i is in a state of stock-out in period t and let $\sigma_{it} = 0$ denote the opposite. Point i incurs a stock-cost of χ_i for all $t \in \mathcal{T}$ where $\sigma_{it} = 1$. For $t \in \mathcal{T}$ where $\sigma_{it} = 1$ and no vehicle $k \in \mathcal{K}$ visits demand point i , an emergency delivery recourse action is applied with a cost of ζ_i . We apply a limited back-order policy where a delivery must be performed in the same period t in which a stock-out occurs. We can limit the probability of stock-out at the demand points to a maximum allowable level γ^{DP} . Demand stochasticity and the probability of stock-out are further discussed in Sections 4.3.1 and 4.3.2.

There is a maximum of one delivery to each demand point per period and this delivery follows a discretized ML policy, which is more general than the OU policy, but less general than the classical ML policy. In the discretized ML policy, the delivery quantity is still part of the decisions, but is chosen from a discrete set as shown in Figure 4.1. Let the set \mathcal{L}_i define for each demand point i its allowable discrete inventory levels. For the case where $\mathcal{L}_i = \{\omega_i\}, \forall i \in \mathcal{P}$, the discretized ML policy reduces to the OU policy. The use of a discretized

Figure 4.1: Discrete Maximum Level Policy Example



ML policy is for the sake of tractability and is further discussed in Section 4.3.2.

Unlike for demand points, our framework ignores inventory tracking at the supply points. In most cases supply point inventories are easier to monitor and manage. And in many situations, for example waste collection (see Chapter 3), tracking supply point inventories is irrelevant. However, in our case this modeling choice is due to the complex propagations of uncertainty that tracking supply point inventory entails. These include but are not limited to 1) the effect of emergency deliveries on the supply point inventories, where it is unclear which supply points will be affected and by how much, and 2) the effect of undelivered quantity on the vehicle when reaching a supply point due to lower than expected demands of the previously visited demand points. Evaluating the above effects is beyond the scope of this research.

A tour executed by vehicle k in period t starts from an origin $o' \in \mathcal{O}'_{kt}$ and ends at a destination $o'' \in \mathcal{O}''_{kt}$ and is a sequence of demand and supply point visits. The maximum tour duration of vehicle k in period t is denoted by H_{kt} . If $H_{kt} = 0$, vehicle k is not available in period t . A tour's origin and destination need not coincide, and the correct definition of the sets \mathcal{O}'_{kt} and \mathcal{O}''_{kt} implies that $\mathcal{O}''_{kt} \cap \mathcal{O}'_{k(t+1)} \neq \emptyset$, i.e there is at least one depot where vehicle k can end its tour in period t and start its tour in period $t + 1$. The correct definition of the above sets also implies that when $H_{kt} = 0$, $\exists o' \in \mathcal{O}'_{kt}$ and $o'' \in \mathcal{O}''_{kt}$ s.t. $\pi_{o'o''} = 0$, i.e there is at least one physical depot at which vehicle k can idle in period t . A penalty Θ is applied on the difference between the minimum and maximum vehicle workload, the latter represented by the total duration of all tours a vehicle executes over the planning horizon. Thus, the penalty serves as an incentive to balance workload among the vehicles.

We distinguish a tour from a trip, the latter being a sequence of demand points \mathcal{S} visited by vehicle k between two supply point visits. The supply point visits delimiting the trips may be in the same or in different periods. In a given solution, the set of supply point delimited trips performed by vehicle k is denoted by \mathfrak{S}_k . Demand stochasticity affects trips through the probability of route failure, which is the probability of the total demand in trip $\mathcal{S} \in \mathfrak{S}_k$ exceeding the vehicle capacity Ω_k . The recourse action is a visit to a supply point. The cost of this recourse action is $C_{\mathcal{S}}$, which is the average routing cost of going from the demand points in \mathcal{S} to their nearest supply point and back. To control its degree of conservatism,

Chapter 4. A Unified Framework for Rich Routing Problems with Stochastic Demands

this cost can be pre-multiplied by a Route Failure Cost Multiplier (RFCM) of ψ . We can also limit the probability of route failure to a maximum allowable level γ^{RF} . The probability of route failure is further discussed in Section 4.3.3. All sets and parameters discussed above are summarized in Table 4.1.

Table 4.1: Notations

Sets			
\mathcal{T}	planning horizon = $\{0, \dots, u\}$	\mathcal{T}^+	shifted planning horizon = $\{1, \dots, u, u+1\}$
\mathcal{O}'_{kt}	set of origins for vehicle k in period t	\mathcal{O}''_{kt}	set of destinations for vehicle k in period t
\mathcal{P}	set of demand points	\mathcal{D}	set of supply points
\mathcal{N}'_{kt}	$= \mathcal{O}'_{kt} \cup \mathcal{O}''_{kt} \cup \mathcal{P} \cup \mathcal{D}$	\mathcal{K}	set of vehicles
\mathcal{C}_i	set of visit period comb. for demand point i	\mathcal{L}_i	set of discrete levels for demand point i
\mathfrak{S}_k	set of trips executed by vehicle k	\mathcal{S}	a particular trip in \mathfrak{S}_k
\mathcal{S}_t	set of demand points in trip \mathcal{S} visited in period t		
Parameters			
φ_k	per-period deployment cost of vehicle k (monetary)		
β_k	unit-distance running cost of vehicle k (monetary)		
θ_k	unit-time running cost of vehicle k (monetary)		
Ω_k	capacity of vehicle k		
π_{ij}	length of arc (i, j)		
$\tau_{ijk t}$	travel time of vehicle k on arc (i, j) in period t		
λ_i, μ_i	lower and upper time window bound at point i		
δ_i	service duration at point i		
ω_i	inventory capacity of demand point i		
ξ_i	visit cost to demand point i (monetary)		
η_i	inventory holding cost at demand point i (monetary)		
ν_i	minimum number of times that demand point i must be visited over the planning horizon		
α_{rt}	1 if period t belongs to visit period combination r , 0 otherwise		
α_{ikt}	1 if point i is accessible by vehicle k in period t , 0 otherwise		
ρ_{it}	stochastic demand of point i in period t		
ε_{it}	stochastic error term of demand point i in period t		
σ_{it}	1 if demand point i is in a state of stock-out in period t , 0 otherwise		
χ_i	stock-out cost at demand point i (monetary)		
ζ_i	emergency delivery cost to demand point i (monetary)		
H_{kt}	maximum tour duration for vehicle k in period t		
Θ	penalty on the difference between the min and max vehicle workload over the planning horizon (monetary)		
ψ	Route Failure Cost Multiplier (RFCM) $\in [0, 1]$		
$C_{\mathcal{S}}$	the average routing cost of going from $\mathcal{S} \in \mathfrak{S}_k$ to the nearest supply point and back to \mathcal{S} (monetary)		
γ^{DP}	maximum allowable probability of stock-out at the demand point in the range of $(0, 1]$		
γ^{RF}	maximum allowable probability of route failure in the range of $(0, 1]$		
Decision Variables			
$x_{ijk t}$	1 if vehicle k traverses arc (i, j) in period t , 0 otherwise (binary)		
y_{ikt}	1 if point i is visited by vehicle k in period t , 0 otherwise (binary)		
z_{kt}	1 if vehicle k is used in period t , 0 otherwise (binary)		
c_{ir}	1 if visit period combination r is assigned to demand point i , 0 otherwise (binary)		
ℓ_{irt}	1 if discrete level r is chosen for demand point i in period t , 0 otherwise (binary)		
q_{ikt}	expected delivery quantity to demand point i by vehicle k in period t (continuous)		
Q_{ikt}	expected cumulative quantity delivered by vehicle k arriving at point i in period t (continuous)		
I_{it}	expected inventory at demand point i at the start of period t (continuous)		
S_{ikt}	start-of-service time of vehicle k at point i in period t (continuous)		
$\underline{b}_{kt}, \bar{b}_{kt}$	lower and upper bound on the tour duration of vehicle k in period t (continuous)		
\underline{B}, \bar{B}	lower and upper bound on the workload for each vehicle (continuous)		

4.3 Capturing Demand Stochasticity

Our framework considers stochastic demands with all other parameters being fully deterministic. Below, we describe in detail how the unified framework captures stochastic demands. In particular, Section 4.3.1 outlines the forecasting of future demands and the minimum amount of forecasting information that the framework needs. Then, Sections 4.3.2 and 4.3.3 derive the probabilities of stock-out and route failure, respectively. We focus specifically on the issue of tractability and the fact that all probability information can be pre-computed or at least partially pre-processed.

4.3.1 Demand Decomposition and Forecasting

We generalize the representation of the stochastic demand introduced in Section 3.2.1 of Chapter 3. Given a demand point $i \in \mathcal{P}$ and a period $t \in \mathcal{T}$, the stochastic demand ρ_{it} decomposes as:

$$\rho_{it} = \mathbb{E}(\rho_{it}) + \varepsilon_{it}, \quad (4.1)$$

where $\mathbb{E}(\rho_{it})$ is the expected demand and ε_{it} is the error component. Let us represent $\varepsilon_{it}, \forall t \in \mathcal{T}, i \in \mathcal{P}$ in the form of a vector as follows:

$$\boldsymbol{\varepsilon} = (\varepsilon_{11}, \dots, \varepsilon_{1|\mathcal{T}|}, \varepsilon_{21}, \dots, \varepsilon_{|\mathcal{P}||\mathcal{T}|}). \quad (4.2)$$

The associated joint distribution is Φ , and $\boldsymbol{\varepsilon} \sim \Phi$ satisfies $\text{var}(\boldsymbol{\varepsilon}) = \mathbf{K}$, with \mathbf{K} representing any covariance structure.

Definition 4.1. *A forecasting model provides the expected demands $\mathbb{E}(\rho_{it}), \forall t \in \mathcal{T}, i \in \mathcal{P}$ and the distribution Φ of $\boldsymbol{\varepsilon}$.*

Any forecasting model that complies with Definition 4.1 can be used. Moreover, the distribution Φ need not be theoretical. The only requirement is that we should be able to simulate it. Therefore, an empirical distribution is also admissible as we can sample from it. The forecasting model thus remains as general as possible, giving freedom for the use of methodologies suitable to the specific application area.

4.3.2 Demand Point Probabilities

Extending the terminology introduced in Section 4.2, and as counterpart to that of Chapter 3, we distinguish between a regular and an emergency delivery to demand point $i \in \mathcal{P}$. Let the binary decision variable $y_{ikt} = 1$ denote a visit to demand point i by vehicle $k \in \mathcal{K}$ in period $t \in \mathcal{T}$, and let $y_{ikt} = 0$ denote otherwise. In other words, a regular delivery to demand point i in period t is one for which $y_{ikt} = 1$ for some vehicle $k \in \mathcal{K}$. Contrarily, an emergency delivery is a recourse action performed in a state of stock-out in the absence of a regular

Chapter 4. A Unified Framework for Rich Routing Problems with Stochastic Demands

delivery, i.e. for $t \in \mathcal{T}$ where $\sigma_{it} = 1$ and $y_{ikt} = 0, \forall k \in \mathcal{K}$. Moreover, an emergency delivery always brings the inventory level at demand point i to its capacity ω_i . That is, for emergency deliveries we restrict the inventory policy to OU. This is in view of preserving tractability and is discussed in further detail below.

To formalize the discussion below, we also introduce the decision variable I_{it} , which represents the expected inventory level of demand point i at the start of period t , and the decision variable q_{ikt} , which represents the expected delivery quantity to demand point i by vehicle k in period t . Using these, we can establish the inventory of point i after delivery in period t as:

$$\Lambda_{it} = I_{it} + \sum_{k \in \mathcal{K}} q_{ikt}. \quad (4.3)$$

Therefore, if $q_{ikt} = 0, \forall k \in \mathcal{K}$, it follows that $\Lambda_{it} = I_{it}$. The two definitions that follow illustrate the information availability over the planning horizon \mathcal{T} and the sequence of actions in each period $t \in \mathcal{T}$.

Definition 4.2. *The initial inventory I_{i0} for each demand point $i \in \mathcal{P}$ is observed and known with certainty. It can be positive, zero or negative.*

As a consequence, the probability of stock-out of any demand point in period $t = 0$ is either 0 or 1.

Definition 4.3. *For each demand point $i \in \mathcal{P}$ and period $t \in \mathcal{T}$, we have: 1) a potential regular delivery which sets Λ_{it} 2) followed by a realization of the demand ρ_{it} . In other words, for a given period t , deliveries take place before demand realizations.*

Given that both the stock-out cost χ_i and the emergency delivery cost ζ_i for demand point i are only paid in a state of stock-out, we are interested in calculating the probability of stock-out for all $i \in \mathcal{P}$ over the planning horizon. To do this, we adapt the ideas presented in Section 3.2.2 of Chapter 3 in the context of the waste collection IRP. Consider a regular delivery to demand point i in period $g \in \mathcal{T}$. We identify four possible ways of reaching a state of stock-out. Given the stochastic demand decomposition formula (4.1) and the action sequence in Definition 4.3, they and their associated probabilities are formulated as:

- Reaching a state of stock-out in period $g + 1$ from a regular delivery in period g . Its probability is unconditional and is given by:

$$\mathbb{P}(\Lambda_{ig} - \rho_{ig} \leq 0) = \mathbb{P}(\varepsilon_{ig} \geq \Lambda_{ig} - \mathbb{E}(\rho_{ig})). \quad (4.4)$$

- Reaching a state of stock-out in periods later than $g + 1$ from a regular delivery in

period g . Its probability is conditional and is given by:

$$\begin{aligned} & \mathbb{P}\left(\Lambda_{ig} - \sum_{t=g}^h \rho_{it} \leq 0 \mid \Lambda_{ig} - \sum_{t=g}^{h-1} \rho_{it} > 0\right) = \\ & = \mathbb{P}\left(\sum_{t=g}^h \varepsilon_{it} \geq \Lambda_{ig} - \sum_{t=g}^h \mathbb{E}(\rho_{it}) \mid \sum_{t=g}^{h-1} \varepsilon_{it} < \Lambda_{ig} - \sum_{t=g}^{h-1} \mathbb{E}(\rho_{it})\right), \quad \forall h > g. \end{aligned} \quad (4.5)$$

- Reaching a state of stock-out in period $g' + 1$ from a state of stock-out in period $g' > g$. Its probability is unconditional and is calculated as a special case of formula (4.4) as follows:

$$\mathbb{P}(\omega_i - \rho_{ig'} \leq 0) = \mathbb{P}(\varepsilon_{ig'} \geq \omega_i - \mathbb{E}(\rho_{ig'})), \quad \forall g' > g. \quad (4.6)$$

- Reaching a state of stock-out in periods later than $g' + 1$ from a state of stock-out in period $g' > g$. Its probability is conditional and is calculated as a special case of formula (4.5) as follows:

$$\begin{aligned} & \mathbb{P}\left(\omega_i - \sum_{t=g'}^h \rho_{it} \leq 0 \mid \omega_i - \sum_{t=g'}^{h-1} \rho_{it} > 0\right) = \\ & = \mathbb{P}\left(\sum_{t=g'}^h \varepsilon_{it} \geq \omega_i - \sum_{t=g'}^h \mathbb{E}(\rho_{it}) \mid \sum_{t=g'}^{h-1} \varepsilon_{it} < \omega_i - \sum_{t=g'}^{h-1} \mathbb{E}(\rho_{it})\right), \quad \forall h > g' > g. \end{aligned} \quad (4.7)$$

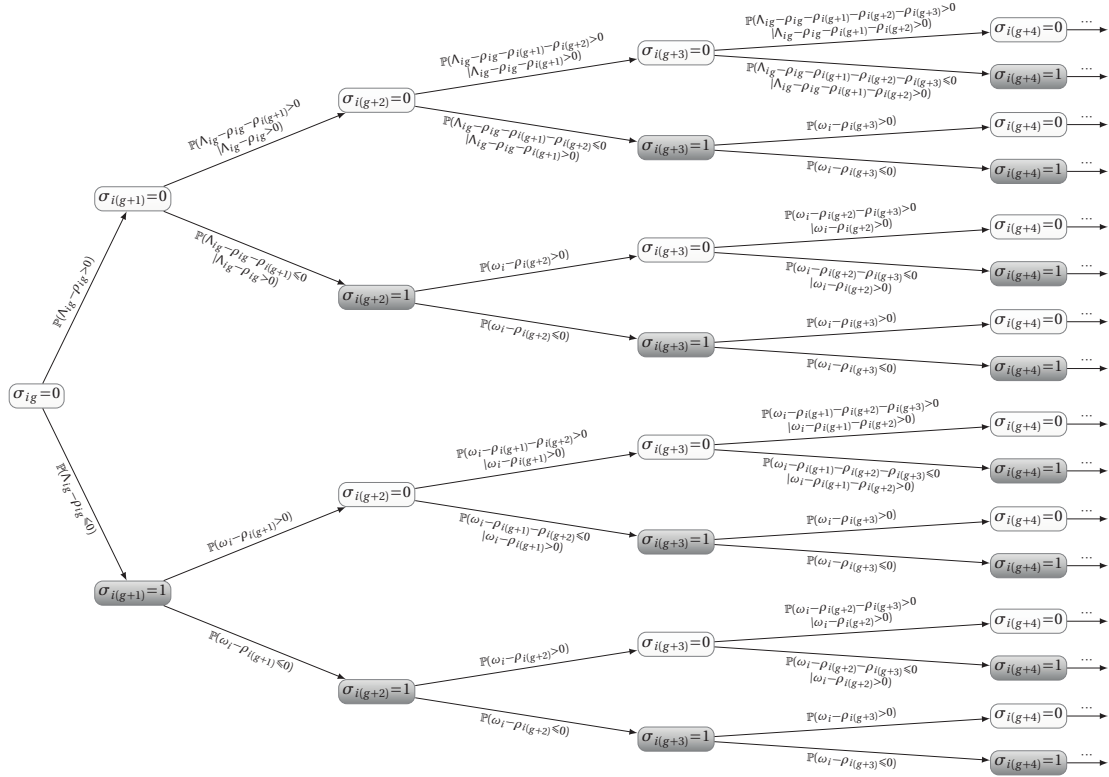
Appendix B proves that the calculation of the probabilities of overflow for a collection problem (see Section 3.2.2 of Chapter 3) is identical. For a demand point i with a regular delivery in period g , the above probabilities are mapped on a binary tree as illustrated in Figure 4.2, in which the state of stock-out is shaded in gray. The probability of stock-out in period $t > g$ is the sum of the probabilities of all possible paths reaching a state $\sigma_{it} = 1$ starting from the root node with an inventory after delivery of Λ_{ig} in period g . The probability of stock-out in period g is calculated on the basis of the previous tree, and is 0 or 1 for $g = 0$. Thus, we arrive at the general expression for the probability of stock-out of demand point i in period t :

$$p_{it}^{\text{DP}} = \mathbb{P}(\sigma_{it} = 1 \mid \Lambda_{im}: m = \max(0, g \in \mathcal{T}: g < t: \exists k \in \mathcal{K}: y_{ikg} = 1)). \quad (4.8)$$

It correctly defines the probability of stock-out as conditional on the inventory after delivery of the most recent regular delivery, identified for each demand point i by the index m . The max operator returns the period 0 if the demand point has not had any regular deliveries prior to period t . For a general distribution Φ , expression (4.8) is non-linear in m and Λ_{im} . Given the discretized ML policy introduced in Section 4.2, we prove the following:

Proposition 4.1. *Under a discretized ML policy, the stock-out probabilities in expression (4.8) can be pre-computed. Moreover, the number of probabilities to pre-compute grows linearly with the number of discrete levels.*

Figure 4.2: Demand Point State Probability Tree



Proof. For the unconditional probabilities (4.4) and (4.6), the number of distinct expressions to evaluate is linear in the number of periods $t \in \mathcal{T}$, while for the conditional probabilities (4.5) and (4.7) it is polynomial. As a consequence, the resulting stock-out probabilities in formula (4.8) can be efficiently pre-computed. Secondly, the formula defines the probability of $\sigma_{it} = 1$ as conditional only on the inventory level Λ_{im} chosen in the most recent delivery period m . The probabilities (4.8) are precomputed for each $r \in \mathcal{L}_i, \forall i \in \mathcal{P}$, hence the number grows linearly with the number of discrete levels. \square

The emergency deliveries still apply an OU policy, otherwise the combinatorial dimension would become intractable. Appendix C.1 demonstrates the use of simulation to pre-compute the stock-out probabilities (4.4)–(4.7) given a general distribution Φ and any covariance structure \mathbf{K} among the error terms $\boldsymbol{\epsilon}$ in formula (4.2). In Section 4.2, it was mentioned that the discretized ML inventory policy is used for the sake of tractability in order to avoid cumbersome calculations at runtime. Indeed, as mentioned in the proof to Proposition 4.1, the ability to pre-compute the stock-out probabilities relies on the discrete values of Λ_{it} .

Finally, for expression (4.8) to be rigorously defined, the value of Λ_{im} must be the expected one. This condition always holds for an OU policy as we deliver up to capacity. However, the ML policy implies a non-negative probability of the chosen Λ_{im} being lower than the realized inventory. There are several possibilities of handling this, including:

- Performing no delivery. This approach leads to a significant increase in complexity due to the additional conditionality now associated with the period m of the most recent regular delivery in formula (4.8), which makes it unattractive.
- Picking up the excess inventory. This approach destroys the monotonicity of the cumulative quantity on the vehicle at each point, increasing the complexity of calculating the probability of route failure, and thus leading to tractability issues (see Section 4.3.3 below).
- Discarding the excess inventory. This approach is the most appealing from a modeling point of view as it allows the use of the expected values Λ_{im} both in the calculation of stock-out and route failure probabilities. Discarding excess inventory can in principle be penalized, its probability being a straightforward extension of formula (4.8). We thus impose the following:

Assumption 4.1. *A regular delivery to demand point $i \in \mathcal{P}$ in period $t \in \mathcal{T}$ discards any inventory above the chosen level Λ_{it} . Thus, a regular delivery sets Λ_{it} according to expectation.*

Assumption 4.1 underlies the calculation of the stock-out probabilities as defined by formula (4.8) as well as the calculation of the route failure probabilities discussed in Section 4.3.3 next.

4.3.3 Route Failure Probabilities

Recalling the notation introduced in Section 4.2, for each vehicle k in a given solution, we identify the set of supply point delimited trips \mathfrak{S}_k . Let the binary decision variables $x_{ijk_t} = 1$ if vehicle k traverses arc (i, j) in period t , and 0 otherwise. For a vehicle k , given $x_{ijk_t}, \forall i, j \in \mathcal{P}, t \in \mathcal{T}$, Algorithm 4.1 builds the set of supply point delimited trips \mathfrak{S}_k , where as before \mathcal{S} is a trip in \mathfrak{S}_k . The algorithm identifies the sequence of visits using the routing variables x_{ijk_t} for each period $t \in \mathcal{T}$. A visit to a supply point starts a new trip \mathcal{S} . In the context of multi-period trips, the supply points delimiting the trips $\mathcal{S} \in \mathfrak{S}_k$ may be visited in different periods t . Thus, each trip \mathcal{S} is further decomposed into sets \mathcal{S}_t , where $\mathcal{S}_t \in \mathcal{S}$ is the set of demand points in trip \mathcal{S} that are visited in period t .

The above notation is used in the formulation of the probability of route failure, which is the probability of the total demand in trip $\mathcal{S} \in \mathfrak{S}_k$ exceeding the vehicle capacity Ω_k . We define the quantity $\Gamma_{\mathcal{S}}$ delivered in trip \mathcal{S} as:

$$\Gamma_{\mathcal{S}} = \sum_{\mathcal{S}_0 \in \mathcal{S}} \sum_{s \in \mathcal{S}_0} (\Lambda_{s0} - I_{s0}) + \sum_{t \in \mathcal{T} \setminus \{0\}} \sum_{\mathcal{S}_t \in \mathcal{S}} \sum_{s \in \mathcal{S}_t} \left(\Lambda_{st} - \Lambda_{sm} + \sum_{h=m}^{t-1} \rho_{sh} \right), \quad (4.9)$$

where $m = \max(0, g \in \mathcal{T} : g < t : \exists k' \in \mathcal{K} : y_{sk'g} = 1)$.

The first summand in formula (4.9) represents the quantity delivered in period $t = 0$ for which there is no uncertainty, while the second summand defines the quantity delivered in

Algorithm 4.1: Construction of the Set of Supply Point Delimited Trips \mathfrak{S}_k for Vehicle k

Input any solution with values of $x_{ijk t}, \forall i, j \in \mathcal{P}, t \in \mathcal{T}$ for vehicle k

Output set of supply point delimited trips \mathfrak{S}_k for vehicle k

```

1:  $\mathcal{S} \leftarrow \mathfrak{S}_k \leftarrow \emptyset$ 
2: for  $t \in \mathcal{T}$  do
3:    $S_t \leftarrow \emptyset$ 
4:    $c \leftarrow j : \sum_{o' \in \mathcal{O}'_{kt}} \sum_{j \in \mathcal{N}_{kt}} x_{o'jkt} = 1$ 
5:   while  $c \notin \mathcal{O}''_{kt}$  do
6:     if  $c \in \mathcal{D}$  then
7:       add  $S_t$  as an element of  $\mathcal{S}$ ; add  $\mathcal{S}$  as an element of  $\mathfrak{S}_k$ 
8:        $S_t \leftarrow \mathcal{S} \leftarrow \emptyset$ 
9:     else if  $c \in \mathcal{P}$  then
10:      add  $c$  as an element of  $S_t$ 
11:    end if
12:     $c \leftarrow j : \sum_{j \in \mathcal{N}_{kt}} x_{c jkt} = 1$ 
13:  end while
14:  add  $S_t$  as an element of  $\mathcal{S}$ 
15: end for

```

periods $t > 0$ given the action sequence in Definition 4.3 and the expected inventory after delivery under Assumption 4.1. Similar to formula (4.8), the index m identifies the most recent regular delivery to point s . Having defined $\Gamma_{\mathcal{S}}$, the probability of route failure in trip $\mathcal{S} \in \mathfrak{S}_k$ performed by vehicle $k \in \mathcal{K}$ becomes:

$$p_{\mathcal{S},k}^{\text{RF}} = \mathbb{P}(\Gamma_{\mathcal{S}} > \Omega_k). \quad (4.10)$$

Formula (4.10) captures the probability of multiple route failures in each trip \mathcal{S} . Unlike in the case of the stock-out probabilities, the probabilities of route failure depend on the optimization decisions, in particular the sets $\mathfrak{S}_k, \forall k \in \mathcal{K}$ at each solution and the values of Λ_{st} and Λ_{sm} . As a consequence, these probabilities cannot be precomputed. Moreover, the distribution of $\Gamma_{\mathcal{S}}$ is generally unknown. The resulting complexity in calculating the route failure probabilities and the need for tractability necessitate the two assumptions below.

Assumption 4.2. *The calculation of the route failure probabilities assumes independent and identically distributed (iid) error terms ε_{it} drawn from any distribution Φ of ε . Consider ε as defined by equation (4.2) above. We impose the iid assumption on the error terms by setting:*

$$\Phi(\varepsilon) = \prod_{t \in \mathcal{T}} \prod_{i \in \mathcal{P}} \Phi'(\varepsilon_{it}), \quad (4.11)$$

where Φ' is the marginal cumulative distribution function of ε_{it} .

Assumption 4.2 is widely used in the literature (Gendreau et al., 2016). In our framework it is only imposed for the calculation of the route failure probabilities, and the distribution Φ is still kept general. Assumption 4.2 renders the demands ρ_{sh} in formula (4.9) independent

of $s \in \mathcal{P}$ and $h \in \mathcal{T}$. Now the distribution of $\Gamma_{\mathcal{S}}$ depends only on the number n of summed demands, which given the action sequence of Definition 4.3 is bounded by $|\mathcal{P}|(|\mathcal{T}| - 1)$. Therefore, an empirical distribution function can be derived for each n and used at runtime. The use of simulation for this partial pre-processing of the route failure probabilities through the derivation of empirical distribution functions is elaborated in Appendix C.2. The numerical experiments in Section 4.6.2 show that for normally distributed error terms the effect on computation time of this approach, in comparison to analytical approximation, is insignificant.

Assumption 4.3. *The calculation of the route failure probabilities ignores the effect of demand points stocking out earlier than expected.*

Including the probability of demand points stocking out earlier than expected leads to complex conditionality given the multiplicity of demand point visits in each trip \mathcal{S} . In addition, it precludes the possibility of partially pre-processing the route failure probabilities as discussed above. The effect of Assumption 4.3 on the objective function is further discussed in Section 4.4.1 and shown to be marginal in the numerical experiments in Section 4.6.2.

4.4 Optimization Model

In Sections 4.4.1, 4.4.2 and 4.4.3 below, we present the objective function and the constraints of the optimization model. It integrates demand stochasticity through the derivation of the probabilities of stock-out and route failure demonstrated in Sections 4.3.2 and 4.3.3, respectively, given all the accompanying definitions and assumptions. The formulations below are presented and interpreted from a distribution point of view. However, since collection can be viewed as the distribution of empty space, the optimization model itself does not change.

To complete the notation, we provide the list of decision variables, including those already used in Section 4.3. Starting with the binary variables, $x_{ijk_t} = 1$ if vehicle k traverses arc (i, j) in period t , 0 otherwise; $y_{ikt} = 1$ if point $i \in \mathcal{O} \cup \mathcal{P} \cup \mathcal{D}$ is visited by vehicle k in period t , 0 otherwise; $z_{kt} = 1$ if vehicle k is used in period t , 0 otherwise; $c_{ir} = 1$ if visit day combination $r \in \mathcal{C}_i$ is assigned to demand point i , 0 otherwise; $\ell_{ir_t} = 1$ if inventory level $r \in \mathcal{L}_i$ is chosen for demand point i in period t , 0 otherwise. Moving to the continuous variables, q_{ikt} is the expected delivery quantity to demand point i by vehicle k in period t ; Q_{ikt} is the expected quantity on vehicle k arriving at point $i \in \mathcal{O} \cup \mathcal{P} \cup \mathcal{D}$ in period t ; I_{it} is the expected inventory at demand point i at the start of period t ; S_{ikt} is the start-of-service time of vehicle k at point $i \in \mathcal{O} \cup \mathcal{P} \cup \mathcal{D}$ in period t ; \underline{b}_{kt} and \bar{b}_{kt} are the lower and upper bound on the tour duration of vehicle k in period t ; and \underline{B} and \bar{B} are the lower and upper bound on the workload for each vehicle. These definitions appear in Table 4.1.

4.4.1 Objective Function

The objective function consists of four deterministic and two stochastic components, all of which are independent of one another. Different combinations of these make it possible to model a variety of routing problems, whether with deterministic and stochastic demands. Starting with the deterministic components, the Expected Inventory Holding Cost (EIHC) is the cost due to keeping the expected inventory at the demand points. Since the inventories in the first period after the end of the planning horizon are completely determined by the decisions taken during the planning horizon, the EIHC is computed for $t \in \mathcal{T} \cup \mathcal{T}^+$, where \mathcal{T}^+ is the planning horizon shifted right by one period. The EIHC is formulated as:

$$\text{EIHC} = \sum_{t \in \mathcal{T} \cup \mathcal{T}^+} \sum_{i \in \mathcal{P}} \eta_i I_{it}. \quad (4.12)$$

The Visit Cost (VC) component applies a cost for each visit to a demand point:

$$\text{VC} = \sum_{t \in \mathcal{T}} \sum_{k \in \mathcal{K}} \sum_{i \in \mathcal{P}} \xi_i y_{ikt}. \quad (4.13)$$

The Routing Cost (RC) component applies the three vehicle-related costs, namely the per-period deployment cost φ_k , the unit-distance running cost β_k and the unit-time running cost θ_k , for each period $t \in \mathcal{T}$ and each vehicle $k \in \mathcal{K}$:

$$\text{RC} = \sum_{t \in \mathcal{T}} \sum_{k \in \mathcal{K}} \left(\varphi_k z_{kt} + \beta_k \sum_{i \in \mathcal{N}_{kt}} \sum_{j \in \mathcal{N}_{kt}} \pi_{ij} x_{ijk t} + \theta_k \left(\sum_{o'' \in \mathcal{O}''_{kt}} S_{o''kt} - \sum_{o' \in \mathcal{O}'_{kt}} S_{o'kt} \right) \right). \quad (4.14)$$

The Workload Balancing (WB) component attempts to balance the workload over the planning horizon equally among the vehicles by penalizing the difference between the lowest and the highest vehicle workload:

$$\text{WB} = \Theta(\bar{B} - \underline{B}). \quad (4.15)$$

Moving to the stochastic components, the Expected Stock-Out and Emergency Delivery Cost (ESOEDC) component, as its name suggests, reflects the stock-out and emergency delivery cost and writes as:

$$\text{ESOEDC} = \sum_{t \in \mathcal{T} \cup \mathcal{T}^+} \sum_{i \in \mathcal{P}} \left(\chi_i + \zeta_i - \zeta_i \sum_{k \in \mathcal{K}} y_{ikt} \right) p_{it}^{\text{DP}}, \quad (4.16)$$

where the probability of stock-out at the demand point p_{it}^{DP} is defined by formula (4.8). For demand point i in period t , the ESOEDC component applies the stock-out cost χ_i and the emergency delivery cost ζ_i in case there is no regular delivery in that period, and only the stock-out cost χ_i in case there is a regular delivery. Although there is no uncertainty in period $t = 0$, we still need to pay the stock-out cost if the demand point is in a state of stock-out.

Since the stock-out probabilities in the first period after the end of the planning horizon are completely determined by the decisions taken during the planning horizon, the ESOEDC is also computed for $t \in \mathcal{T} \cup \mathcal{T}^+$.

The Expected Route Failure Cost (ERFC) captures the risk of the vehicles running out of capacity before reaching the next scheduled visit to a supply point due to higher than expected demands. It is expressed as:

$$\text{ERFC} = \sum_{k \in \mathcal{K}} \sum_{\mathcal{S} \in \mathfrak{S}_k} \psi C_{\mathcal{S}} p_{\mathcal{S},k}^{\text{RF}}, \quad (4.17)$$

where, as in Section 4.3.3, \mathfrak{S}_k is the set of supply point delimited trips executed by vehicle k , $\mathcal{S} \in \mathfrak{S}_k$ is a particular trip in that set, and $C_{\mathcal{S}}$ is the average routing cost of going from the demand points in \mathcal{S} to the nearest supply point and back. The parameter $\psi \in [0, 1]$, which we refer to as the Route Failure Cost Multiplier (RFCM), is used to scale up or down the degree of conservatism of the ERFC component.

The resulting objective function z is non-linear due to the non-linear nature of the ESOEDC and ERFC components. In its general form, it is formulated as:

$$\min z = \text{EIHC} + \text{VC} + \text{RC} + \text{WB} + \text{ESOEDC} + \text{ERFC}. \quad (4.18)$$

The RC, ESOEDC and ERFC components are generalized from Section 3.2.2 of Chapter 3. The optimization model can be applied in a rolling horizon fashion, as described in Section 3.4.3 of Chapter 3. That is, the objective function is evaluated over the planning horizon, but the decisions to implement are those in period $t = 0$. As a consequence, the decisions to implement in period $t = 0$ are forward-looking. After they are implemented, the planning horizon is rolled over by one period and the problem is solved again. Thus, at each rollover we include more information about the future.

Objective Function Overestimation of the Real Cost.

While the ESOEDC component captures the probability of demand points stocking out in each period of the planning horizon, the rest of the components do not, as the probability expressions would become intractable (see Section 4.3.3 above). Trudeau and Dror (1992) solve a stochastic IRP with the assumption of a single delivery and stock-out for each demand point over the planning horizon. Given this setup, Trudeau and Dror (1992) come up with analytical expressions of the effect on the routing and route failure cost of demand points stocking out earlier than expected. Given their assumptions, if a demand point stocks out earlier than expected, it is simply skipped in the tours. The generality of our framework prevents us from developing tractable analytical expressions of the effect of demand points stocking out earlier than expected on all components of the objective function. Yet, we can to a certain extent analyze the mismatch between the modeled objective function cost and the real cost. Let us introduce the following:

Chapter 4. A Unified Framework for Rich Routing Problems with Stochastic Demands

Definition 4.4. *Given a scenario with a demand point stocking out earlier than expected, a reaction policy is defined as the adjustment of the subsequent decisions in response to the emergency delivery. It is important to distinguish between the recourse action, this being the emergency delivery, and the reaction policy.*

Reaction policies can vary from doing nothing to completely re-optimizing the subsequent decisions. Given the un-captured effect of demand points stocking out earlier than expected, we prove the following:

Proposition 4.2. *In the absence of the EIHC component, objective function (4.18) overestimates the real cost.*

Proof. Consider demand point $i \in \mathcal{P}$ that stocks out in period g and is not visited for a regular delivery in period g . For a do-nothing reaction policy, there is no effect on the VC, RC and WB components as it implies no change in the routing decisions. The ESOEDC component already captures the probability of demand points stocking out in each period of the planning horizon. For the effect on the ERFC component, we identify two cases:

1. There is a vehicle $k \in \mathcal{K}$ that visits point i for a regular delivery during trip $\mathcal{S} \in \mathfrak{S}_k$ in period $t > g$. Given the emergency delivery to point i in period g , vehicle k will deliver less than expected in trip \mathcal{S} , reducing the probability of route failure $p_{\mathcal{S},k}^{\text{RF}}$ according to formula (4.10).
2. Alternatively, there is no trip \mathcal{S} that visits point i in period $t > g$. Therefore, $p_{\mathcal{S},k}^{\text{RF}}$ remains unaffected for all trips $\mathcal{S} \in \mathfrak{S}_k, \forall k \in \mathcal{K}$.

Given the existence of a more sophisticated reaction policy, with an optimal reaction policy at the extreme, the overestimation of the real cost may be higher. The above discussion assumes out the EIHC component. In a distribution problem, a stock-out in period g , followed by an emergency delivery, results in inventory levels I_{it} being higher than expected for $t > g$. Thus, expression (4.12) of the EIHC underestimates the inventory holding cost for $t > g$. It is the contrary for a collection problem, where an overflow in period g , followed by an emergency collection, results in inventory levels I_{it} being lower than expected for $t > g$. In this case, expression (4.12) of the EIHC overestimates the inventory holding cost. \square

The overestimation due to the do-nothing reaction policy is straightforward to evaluate through simulation on the final solution. Contrarily, the evaluation of the effect of an optimal reaction policy would require the re-optimization of the decisions after each stock-out. We can avoid the computational burden of the latter by computing bound information on the maximum overestimation. This is further explored in the numerical experiments in Section 4.6.2.

4.4.2 Deterministic Constraints

The deterministic constraints are extended from those presented in Section 3.2.2 of Chapter 3. Starting with the basic routing constraints, tours must have an origin and a destination depot, as ensured by constraints (4.19), which also allow for simple relocation tours not visiting any demand or supply points. Constraints (4.20) and (4.21) forbid returns to the origin depots and departures from the destination depots. Given the possibility of open tours, we need to ensure that a vehicle's destination depot in period t is the same as its origin depot in period $t + 1$. Constraints (4.22) propagate this condition through the planning horizon. Further on, constraints (4.23) and (4.24) link the visit and the routing variables, and constraints (4.25) ensure that a demand point is visited at most once per period. Accessibility restrictions and continuity of service are enforced by constraints (4.26). Constraints (4.27) ensure flow conservation.

$$\sum_{o' \in \mathcal{O}'_{kt}} \sum_{j \in \mathcal{N}_{kt}} x_{o'jkt} = \sum_{i \in \mathcal{N}_{kt}} \sum_{o'' \in \mathcal{O}''_{kt}} x_{io''kt}, \quad \forall t \in \mathcal{T}, k \in \mathcal{K} \quad (4.19)$$

$$\sum_{i \in \mathcal{N}_{kt}} x_{io'kt} = 0, \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, o' \in \mathcal{O}'_{kt} \quad (4.20)$$

$$\sum_{j \in \mathcal{N}_{kt}} x_{o''jkt} = 0, \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, o'' \in \mathcal{O}''_{kt} \quad (4.21)$$

$$\sum_{i \in \mathcal{N}_{kt}} x_{iokt} = \sum_{j \in \mathcal{N}_{k(t+1)}} x_{ojk(t+1)}, \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, o \in \mathcal{O}''_{kt} \cap \mathcal{O}'_{k(t+1)} \quad (4.22)$$

$$y_{ikt} = \sum_{j \in \mathcal{N}_{kt}} x_{ijk t}, \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, i \in \mathcal{N}_{kt} \setminus \mathcal{O}''_{kt} \quad (4.23)$$

$$y_{jkt} = \sum_{i \in \mathcal{N}_{kt}} x_{ijk t}, \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, j \in \mathcal{O}''_{kt} \quad (4.24)$$

$$\sum_{k \in \mathcal{K}} y_{ikt} \leq 1, \quad \forall t \in \mathcal{T}, i \in \mathcal{P} \quad (4.25)$$

$$y_{ikt} \leq \alpha_{ikt}, \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, i \in \mathcal{P} \cup \mathcal{D} \quad (4.26)$$

$$\sum_{i \in \mathcal{N}_{kt}} x_{ijk t} = \sum_{i \in \mathcal{N}_{kt}} x_{jik t}, \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, j \in \mathcal{P} \cup \mathcal{D} \quad (4.27)$$

The periodicity aspect is established by constraints (4.28), which assign exactly one visit period combination to each demand point, and constraints (4.29), which in turn limit visits to the periods corresponding to the assigned visit period combination (Cordeau et al., 1997). The set \mathcal{C}_i may contain visit period combinations with different frequencies, which makes the visit frequency part of the optimization decisions.

$$\sum_{r \in \mathcal{C}_i} c_{ir} = 1, \quad \forall i \in \mathcal{P} \quad (4.28)$$

$$\sum_{k \in \mathcal{K}} y_{ikt} - \sum_{r \in \mathcal{C}_i} \alpha_{rt} c_{ir} = 0, \quad \forall t \in \mathcal{T}, i \in \mathcal{P} \quad (4.29)$$

The inventory constraints at the demand points comply with the action sequence in Definition 4.3. Constraints (4.30) track the expected inventory in period t as a function of the expected inventory, the quantity delivered to the point, and its expected demand in period $t - 1$. Constraints (4.31) ensure that the expected inventory remains non-negative, and constraints (4.32) force a delivery if the inventory is below zero in period $t = 0$. Constraints (4.33)–(4.36) define the choice of a discrete inventory level and the delivery quantity it entails. In particular, constraints (4.33) stipulate that if a demand point is visited, then a discrete inventory level after delivery must be chosen. Constraints (4.34) and (4.35) provide a lower and an upper bound on the delivery quantity which, if the point is visited, is equal to the difference between the chosen discrete inventory level after delivery and the expected inventory. The latter also imply that if the point is visited, the chosen level will be higher than the expected inventory. Constraints (4.36) force the delivery quantity to zero if the point is not visited. The big- M values in constraints (4.34) and (4.36) are equal to $2\omega_i$ for $t = 0$ and to ω_i otherwise, reflecting the fact that the expected delivery quantity cannot exceed demand point capacity, except in period $t = 0$.

$$I_{it} = I_{i(t-1)} + \sum_{k \in \mathcal{K}} q_{ik(t-1)} - \mathbb{E}(\rho_{i(t-1)}), \quad \forall t \in \mathcal{T}^+, i \in \mathcal{P} \quad (4.30)$$

$$I_{it} \geq 0, \quad \forall t \in \mathcal{T}^+, i \in \mathcal{P} \quad (4.31)$$

$$-I_{i0} \leq \omega_i \sum_{k \in \mathcal{K}} y_{ik0}, \quad \forall i \in \mathcal{P} \quad (4.32)$$

$$\sum_{k \in \mathcal{K}} y_{ikt} - \sum_{r \in \mathcal{L}_i} \ell_{irt} = 0, \quad \forall t \in \mathcal{T}, i \in \mathcal{P} \quad (4.33)$$

$$q_{ikt} \geq \sum_{r \in \mathcal{L}_i} r \ell_{irt} - I_{it} - M(1 - y_{ikt}), \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, i \in \mathcal{P} \quad (4.34)$$

$$q_{ikt} \leq \sum_{r \in \mathcal{L}_i} r \ell_{irt} - I_{it} + \omega_i(1 - y_{ikt}), \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, i \in \mathcal{P} \quad (4.35)$$

$$q_{ikt} \leq M y_{ikt}, \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, i \in \mathcal{P} \quad (4.36)$$

In the context of vehicle capacities, constraints (4.37) limit the cumulative quantity delivered by the vehicle at each demand point, while constraints (4.38) reset it to zero at the supply points. Keeping track of the cumulative quantity delivered by the vehicle is achieved by constraints (4.39). In the context of multi-period trips, constraints (4.40) link the quantity delivered by the vehicle from one period to the next. Forcing the vehicle to visit a supply point immediately after the origin depot or immediately before the destination depot applies to certain problems and is exemplified in Section 4.5 next.

$$q_{ikt} \leq Q_{ikt} \leq \Omega_k, \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, i \in \mathcal{P} \quad (4.37)$$

$$Q_{ikt} = 0, \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, i \in \mathcal{D} \quad (4.38)$$

$$Q_{ikt} + q_{jkt} \leq Q_{jkt} + \Omega_k(1 - x_{ijkt}), \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, i \in \mathcal{N}_{kt}, j \in \mathcal{N}_{kt} \setminus \mathcal{D} \quad (4.39)$$

$$Q_{o'k(t+1)} \geq Q_{o''kt}, \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, o' \in \mathcal{O}'_{kt}, o'' \in \mathcal{O}''_{kt} \quad (4.40)$$

The next set of constraints expresses the intra-period temporal characteristics of the problem. Constraints (4.41) calculate the start-of-service time at each point and eliminate the possibility of subtours. Constraints (4.42) enforce the time windows. Constraints (4.43) bound the tour duration from above and below. Constraints (4.44) enforce the maximum tour duration, and with it availabilities and vehicle use. Constraints (4.45) and (4.46) bound the total tour duration over the planning horizon for each vehicle. The difference between \underline{B} and \bar{B} is the difference between the lowest and highest vehicle workload over the planning horizon, which is penalized by the WB component in the objective function.

$$S_{ikt} + \delta_i + \tau_{ijkt} \leq S_{jkt} + (\mu_i + \delta_i + \tau_{ijkt})(1 - x_{ijkt}), \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, i \in \mathcal{N}_{kt}, j \in \mathcal{N}_{kt} \quad (4.41)$$

$$\lambda_i y_{ikt} \leq S_{ikt} \leq \mu_i y_{ikt}, \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, i \in \mathcal{N}_{kt} \quad (4.42)$$

$$\underline{b}_{kt} \leq \sum_{o'' \in \mathcal{O}''_{kt}} S_{o''kt} - \sum_{o' \in \mathcal{O}'_{kt}} S_{o'kt} \leq \bar{b}_{kt}, \quad \forall t \in \mathcal{T}, k \in \mathcal{K} \quad (4.43)$$

$$\bar{b}_{kt} \leq H_{kt} z_{kt}, \quad \forall t \in \mathcal{T}, k \in \mathcal{K} \quad (4.44)$$

$$\underline{B} \leq \sum_{t \in \mathcal{T}} \underline{b}_{kt}, \quad \forall k \in \mathcal{K} \quad (4.45)$$

$$\bar{B} \geq \sum_{t \in \mathcal{T}} \bar{b}_{kt}, \quad \forall k \in \mathcal{K} \quad (4.46)$$

Finally, lines (4.47)–(4.48) establish the variable domains.

$$x_{ijkt}, y_{ikt}, z_{kt}, c_{ir'}, \ell_{ir''t} \in \{0, 1\}, \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, i, j \in \mathcal{N}_{kt}, r' \in \mathcal{C}_i, r'' \in \mathcal{L}_i \quad (4.47)$$

$$q_{ikt}, Q_{ikt}, I_{it}, S_{ikt}, \underline{b}_{kt}, \bar{b}_{kt}, \underline{B}, \bar{B} \geq 0, \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, i \in \mathcal{N}_{kt} \quad (4.48)$$

4.4.3 Probabilistic Constraints

As an alternative to integrating stochastic demand information in the objective function through the ESOEDC and the ERFC components, it can be included at the constraint level in the form of probabilistic constraints. Constraints (4.49) and (4.50) below impose a maximum allowable probability of stock-out and route failure, respectively.

$$p_{it}^{\text{DP}} \leq \gamma^{\text{DP}}, \quad \forall t \in \mathcal{T}, i \in \mathcal{P} \quad (4.49)$$

$$p_{\mathcal{S},k}^{\text{RF}} \leq \gamma^{\text{RF}}, \quad \forall k \in \mathcal{K}, \mathcal{S} \in \mathfrak{S}_k \quad (4.50)$$

4.5 Application Examples

The framework developed and presented in Sections 4.2, 4.3 and 4.4 can be applied to problems from different fields of routing and logistics optimization. In the sections below, we discuss in more detail a vehicle routing problem, a health care inventory routing problem, a waste collection inventory routing problem, a maritime inventory routing problem, and a facility maintenance problem.

4.5.1 The Vehicle Routing Problem

In a VRP setting, the presence of stochastic demands may lead to route failures but stock-outs do not apply. To adapt the framework, we define a planning horizon $\mathcal{T} = \{0, 1, 2\}$, s.t. $H_{k0} = H_{k2} = 0, \forall k \in \mathcal{K}$, i.e. the planning horizon consists of three periods and no vehicle is available in periods $t = 0$ and $t = 2$. Moreover, $I_{i0} = \omega_i$ and $\mathcal{L}_i = \{\omega_i\}, \forall i \in \mathcal{P}$, i.e. the initial inventory of all demand points is equal to capacity and we apply an OU inventory policy. Given the action sequence of Definition 4.3, the visits to the demand points deliver the demands ρ_{i0} realized in period 0. The VRP is a single-period problem and the fact that it is effectively solved for period $t = 1$ is of no consequence. In model (VRP), the objective (4.51) consists of the RC and the ERFC components. Given constraints (4.38) and (4.39), constraints (4.52) force a visit to a supply point immediately after the origin depot. Constraints (4.25) are replaced by constraints (4.53) below to enforce a delivery to each demand point in period $t = 1$, a necessary condition for a feasible VRP solution. The periodicity related constraints (4.28) and (4.29) are dropped as they become irrelevant for a single period.

$$(VRP) \quad \min \quad z = RC + ERFC \quad (4.51)$$

$$\text{s.t.} \quad \text{Constraints (4.19)–(4.24), (4.26)–(4.27), (4.30)–(4.48)}$$

$$Q_{o'k1} = \Omega_k, \quad \forall k \in \mathcal{K}, o' \in \mathcal{O}'_{k1} \quad (4.52)$$

$$\sum_{k \in \mathcal{K}} y_{ik1} = 1, \quad \forall i \in \mathcal{P} \quad (4.53)$$

4.5.2 The Health Care Inventory Routing Problem

The health care IRP generalizes the nurse routing and scheduling problem, in which nurses visit patient homes to provide treatment. In this problem, \mathcal{P} is the set of patient homes and \mathcal{D} is the set of medical facilities. In addition to providing treatment, nurses deliver medications with stochastic demand. Continuity of care and workload balancing, which are the two paramount concerns in the nurse routing problem, are supported by the framework. As is the periodic aspect, given that medical treatments typically have to be performed with a certain frequency. Pricing can also be introduced in the setup via a negative visit cost. We keep the model (HCIRP) general, including all constraints, and with the objective function (4.54) including all but the EIHC component.

$$(HCIRP) \quad \min \quad z = VC + RC + WB + ESOEDC + ERFC \quad (4.54)$$

$$\text{s.t.} \quad \text{Constraints (4.19)–(4.48)}$$

4.5.3 The Waste Collection Inventory Routing Problem

This is the problem discussed in detail in Chapter 3. In this IRP variant, trucks collect waste from containers with stochastic demands. Here, \mathcal{P} denotes the set of waste containers and \mathcal{D} denotes the set of disposal facilities. The unified framework formulated in Sections 4.2, 4.3

and 4.4 can be applied with minimal changes by relabeling the problem as the distribution of empty space. The objective function of model (WCIRP) mimics the objective in Chapter 3 and includes the RC, ESOEDC and ERFC components. Given constraints (4.38) and (4.39), constraints (4.56) force a visit to a disposal facility immediately before the destination depot.

$$\text{(WCIRP) } \min \quad z = \text{RC} + \text{ESOEDC} + \text{ERFC} \quad (4.55)$$

$$\text{s.t.} \quad \text{Constraints (4.19)–(4.48)}$$

$$Q_{o''kt} = 0, \quad \forall t \in \mathcal{T}, k \in \mathcal{K}, o'' \in \mathcal{O}''_{kt} \quad (4.56)$$

4.5.4 The Maritime Inventory Routing Problem

In this problem, a fleet of ships transports a commodity from a set \mathcal{D} of supply terminals to a set \mathcal{P} of demand terminals. A particular feature of this application is that emergency deliveries may be impractical due to long shipping distances, which would make the state of stock-out at a demand terminal a final state. This can be achieved simply by setting the probabilities defined by expression (4.6) to one. Since emergency deliveries are not performed, the emergency delivery cost $\zeta_i = 0, \forall i \in \mathcal{P}$. Maritime routing problems are also characterized by open and multi-period tours, which may include idling. In our framework, constraints (4.19) allow for open tours, while multi-period tours are enabled by defining the set of depots so that $\exists o \in \mathcal{O}$ s.t. $\pi_{oi} = \pi_{io} = 0, \forall i \in \mathcal{P} \cup \mathcal{D}$ and $\mathcal{O}'_{kt} \equiv \mathcal{O}''_{kt} \equiv \mathcal{O}, \forall t \in \mathcal{T}, k \in \mathcal{K}$, or in other words there is an origin and a destination depot at zero distance from each demand and supply terminal. A tour can thus effectively end at a demand or supply terminal in period t and start from it in period $t + 1$. This graph extension is a modeling feature that can be efficiently exploited in the solution methodology. The objective function of model (MIRP) includes all but the WB component. The VC component, in particular, may be used to capture terminal docking fees.

$$\text{(MIRP) } \min \quad z = \text{EIHC} + \text{VC} + \text{RC} + \text{ESOEDC} + \text{ERFC} \quad (4.57)$$

$$\text{s.t.} \quad \text{Constraints (4.19)–(4.48)}$$

4.5.5 The Facility Maintenance Problem

The facility maintenance problem is a probability-based routing problem in which a set of facilities is visited by a set of technicians for inspection. In this problem, the set \mathcal{P} represents the facilities, while the set \mathcal{D} is irrelevant. Uncertainty with respect to breakdowns can be considered as accumulating in a fashion similar to that of inventory. Consider facility $i \in \mathcal{P}$ in period t . We can interpret the state $\sigma_{it} = 1$ as a breakdown, and the state $\sigma_{it} = 0$ as operational. If a facility is in a state of breakdown in period t , an emergency visit must be performed to repair it. The probability of breakdown p_{it}^{DP} of facility i in period t is adapted

Chapter 4. A Unified Framework for Rich Routing Problems with Stochastic Demands

from expression (4.8) as a function of the most recent visit to the facility and is modeled as:

$$p_{it}^{\text{DP}} = \mathbb{P}(\sigma_{it} = 1 \mid g \in \mathbb{Z}: g < t: \exists k \in \mathcal{K}: y_{ikg} = 1). \quad (4.58)$$

The use of the set \mathbb{Z} , which includes the negative integers, implies that the most recent visit may be before the start of the planning horizon \mathcal{T} . The states $\sigma_{i0}, \forall i \in \mathcal{P}$ are known with certainty. The objective function (4.59) in model (FMP1) is the sum of routing cost and the Expected Emergency Repair Cost (EERC). All inventory related constraints (4.30)–(4.36) and vehicle capacity related constraints (4.37)–(4.40) are irrelevant and are hence dropped. The new set of constraints (4.60) is added to force a visit to a facility in a state of breakdown.

$$\text{(FMP1) } \min \quad z = \text{RC} + \text{EERC} \quad (4.59)$$

$$\text{s.t.} \quad \text{Constraints (4.19)–(4.29), (4.41)–(4.48)}$$

$$\sum_{k \in \mathcal{K}} y_{ikt} = 1, \quad \forall t \in \mathcal{T}, i \in \mathcal{P}: p_{it}^{\text{DP}} = 1 \quad (4.60)$$

The EERC is a reformulation of the ESOEDC from formula (4.16) and is expressed as:

$$\text{EERC} = \sum_{t \in \mathcal{T} \cup \mathcal{T}^+} \sum_{i \in \mathcal{P}} p_{it}^{\text{DP}} \zeta_i. \quad (4.61)$$

Since the probabilities in the facility maintenance problem are provided exogenously, as opposed to being calculated based on demand stochasticity, an alternative formulation involving the probabilistic constraints (4.49) is given in model (FMP2). Since the treatment of the probability of breakdown is in the constraints, the objective (4.62) is routing-only.

$$\text{(FMP2) } \min \quad z = \text{RC} \quad (4.62)$$

$$\text{s.t.} \quad \text{Constraints (4.19)–(4.29), (4.41)–(4.48)}$$

$$\text{Constraints (4.49)}$$

$$\text{Constraints (4.60)}$$

Given that the facility maintenance problem considers no demands, unlike in the case of the waste collection IRP, there is no deterministic equivalent problem that simply ignores the stochastic components. We could imagine several deterministic policies, for example periodicity-based visits enforced by constraints (4.28)–(4.29). A more flexible deterministic alternative would be visiting a facility $i \in \mathcal{P}$ at least ν_i times over the planning horizon. In the model (FMPD) below, this is ensured by constraints (4.63).

$$\text{(FMPD) } \min \quad z = \text{RC}$$

$$\text{s.t.} \quad \text{Constraints (4.19)–(4.29), (4.41)–(4.48)}$$

$$\text{Constraints (4.60)}$$

$$\sum_{t \in \mathcal{T}} \sum_{k \in \mathcal{K}} y_{ikt} \geq \nu_i, \quad \forall i \in \mathcal{P} \quad (4.63)$$

4.6 Numerical Experiments

In the following, we carry out a series of experiments to investigate various features of the proposed unified framework. Section 4.6.1 introduces the instance sets with the new set of facility maintenance instances. Section 4.6.2 sets the background with the main conclusions of Chapter 3 on the waste collection IRP instances and performs further experiments on this set. In particular, it studies the effect on tractability of using empirical distribution functions for calculating the route failure probabilities at runtime, and analyzes the objective function's overestimation of the real cost. Section 4.6.3 presents the new case study based on the facility maintenance problem. Various solution methodologies may be appropriate for our unified framework, as long as they can handle the probability-based calculations and support its rich routing features. We use the ALNS developed in Chapter 3, which was shown to have excellent performance on VRP and IRP benchmark instances from the literature and to be very stable on the waste collection IRP instances. The ALNS is implemented as a single-thread application in Java and the probability calculations for the state probability trees (Figure 4.2) are performed in R. All tests have been run on a 3.33 GHz Intel Xeon X5680 server running a 64-bit Ubuntu 16.04.2. In all experiments, each instance is solved 10 times.


4.6.1 Instances

The waste collection IRP instances introduced in Section 3.4.3 of Chapter 3 are 63 instances of white glass collections performed in the canton of Geneva, Switzerland in the years 2014, 2015 and 2016. A map of the collection area was presented in Chapter 2. In these instances, demands are forecast using the count data mixture model presented in Section 3.2.1 of Chapter 3 using the previous 90 days of data, and assuming iid normal error terms ε_{it} for all $i \in \mathcal{P}$ and $t \in \mathcal{T}$, which is supported by the data. Absence of historical container level data prevents demand forecasting for certain weeks of the sample period, for which instances are not generated.

The second set consists of 94 instances of the facility maintenance problem with an average of 42 facilities and a maximum of 62. These instances are built from the same data used for building the waste collection IRP instances. However, since the facility maintenance problem described in Section 4.5.5 does not consider demands, we are not limited by the absence of historical container level data. Hence, the 94 instances of the facility maintenance problems vs. the 63 instances of the waste collection IRP. For each facility $i \in \mathcal{P}$, we set a service duration of 30 minutes, and tours are now constrained to a maximum duration of eight hours, instead of four. The probability of breakdown is modeled using the cumulative distribution function of the log-logistic distribution. That is, the probability p_{it}^{DP} of breakdown of facility i in period t defined in formula (4.58) is given by:

$$p_{it}^{\text{DP}} = \frac{1}{1 + \left(\frac{t-g}{a}\right)^{-\beta}}, \quad (4.64)$$

Figure 4.3: Breakdown Probabilities for Different Values of α



where g is the period of the most recent visit. We set the value of β to 5, while α is randomly chosen for each facility as an integer between 10 and 15, inclusive. Figure 4.3 plots the breakdown probability for the different values of the α parameter. The probability accumulates in a way similar to how inventory builds up in the IRP. In addition, for each facility in period 0, we draw a random integer between 1 and 3, inclusive, for the number of days since the most recent visit.

4.6.2 Solving the Waste Collection Inventory Routing Problem

In Chapter 3, we develop an ALNS that exhibits excellent performance on VRP and IRP benchmark instances from the literature. Using it, we demonstrate that our stochastic model performs significantly better compared to alternative deterministic policies in its ability to reduce the occurrence of container overflows for the same routing cost. Here, we conduct further experiments on these instances. In particular, we assess the effect on tractability of using empirical distribution functions at runtime for calculating the route failure probabilities, and analyze the objective function's overestimation of the real cost previously discussed in Section 4.4.1. A note worth mentioning is that the waste collection IRP presented in Chapter 3 considers a single depot, while this chapter extends the setup to multiple depots. In this context, Appendix A.3 evaluates the benefit of allowing open tours with different origin and destination depots, with a vehicle's destination depot on day t becoming its origin depot on day $t + 1$. The results indicate that the benefit of open tours first discussed in Chapter 2 holds and may even be more convincing in a multi-day setting. Appendix A.3 also describes a new depot replacement operator that is added to the list of ALNS repair operators in Section 3.3.2 of Chapter 3.

Assessing the Effect of Empirical Distribution Functions on Tractability.

As described in Section 4.3.3, assuming iid error terms drawn from any distribution Φ allows the partial pre-processing of the route failure probabilities through the derivation of empirical distribution functions to be used at runtime. Clearly, the main risk of using empirical

distribution functions is their impact on tractability and the precision of the resulting probability. To investigate this, we use the simulation methodology described in Appendix C.2 to build Empirical Cumulative Distribution Functions (ECDFs) for $M = 100,000$ draws. The ECDFs are constructed using the `EmpiricalDistribution` class of the Apache Commons Math 3.6.1 release¹. We test two configurations for the ECDFs, one binning the draws in 1000 bins and one binning them in 100 bins. Computational experiments show that the configuration with 1000 bins exhibits a squared error with respect to the normal distribution in the order of 10^{-7} , while for the configuration with 100 bins, it is in the order of 10^{-6} .

Table 4.2 reports the results of the experiments. The experiments are performed for an Emergency Collection Cost (ECC) $\zeta_i = 100$ CHF for all containers $i \in \mathcal{P}$ and a Route Failure Cost Multiplier (RFCM) $\psi = 1$. In the table, each row reports averaged values over the 63 instances. The first column identifies the version of the ALNS used, i.e. the original one of Chapter 3 vs. the one using ECDFs, while the second column identifies the binning configuration. The original ALNS uses the analytical approximation of the normal distribution of Abramowitz and Stegun (1972), which is possible given the normality of the error terms of the waste collection IRP instances. The next two columns show the ECC and the RFCM, which are the same for all instances. The fifth, sixth and seventh columns present the best, average and worst cost over 10 runs. In a similar fashion, the eighth, ninth and tenth columns report the best, average and worst computation time, and the eleventh, twelfth and thirteenth columns report the best, average and worst number of calls to the ECDFs over 10 runs. Expectedly, Table 4.2 shows that the different implementations have no impact on the solution cost. We also observe that the implementation with 100 bins has a computation time that is virtually the same as that of the original implementation. However, as mentioned before, the binning configuration with 1000 bins has a squared error which is one degree of magnitude lower, while its computation time is only about 5% higher. Therefore, this configuration may be preferable. In summary, unless the distribution of the forecasting error terms adheres to the simple convolution property, as in the case of the normal distribution, the route failure probabilities cannot be evaluated analytically at runtime. Nevertheless, the results of Table 4.2 indicate that using pre-processed ECDFs to calculate the route failure probabilities preserves tractability and has a negligible impact on computation time.

Table 4.2: Impact of Empirical Distribution Functions on Tractability

ALNS version	Bins	ECC	RFCM	Cost (CHF)			Runtime (s.)			ECDF calls (millions)		
				Best	Avg	Worst	Best	Avg	Worst	Best	Avg	Worst
Original	–	100.00	1.00	662.65	666.64	672.87	870.65	906.84	936.40	–	–	–
ECDFs	1000	100.00	1.00	662.63	666.74	673.35	909.06	948.77	982.68	52.95	58.90	65.00
ECDFs	100	100.00	1.00	662.49	666.46	672.73	869.52	903.81	932.79	52.94	58.44	63.90

¹<http://commons.apache.org/proper/commons-math/javadocs/api-3.6.1/index.html>

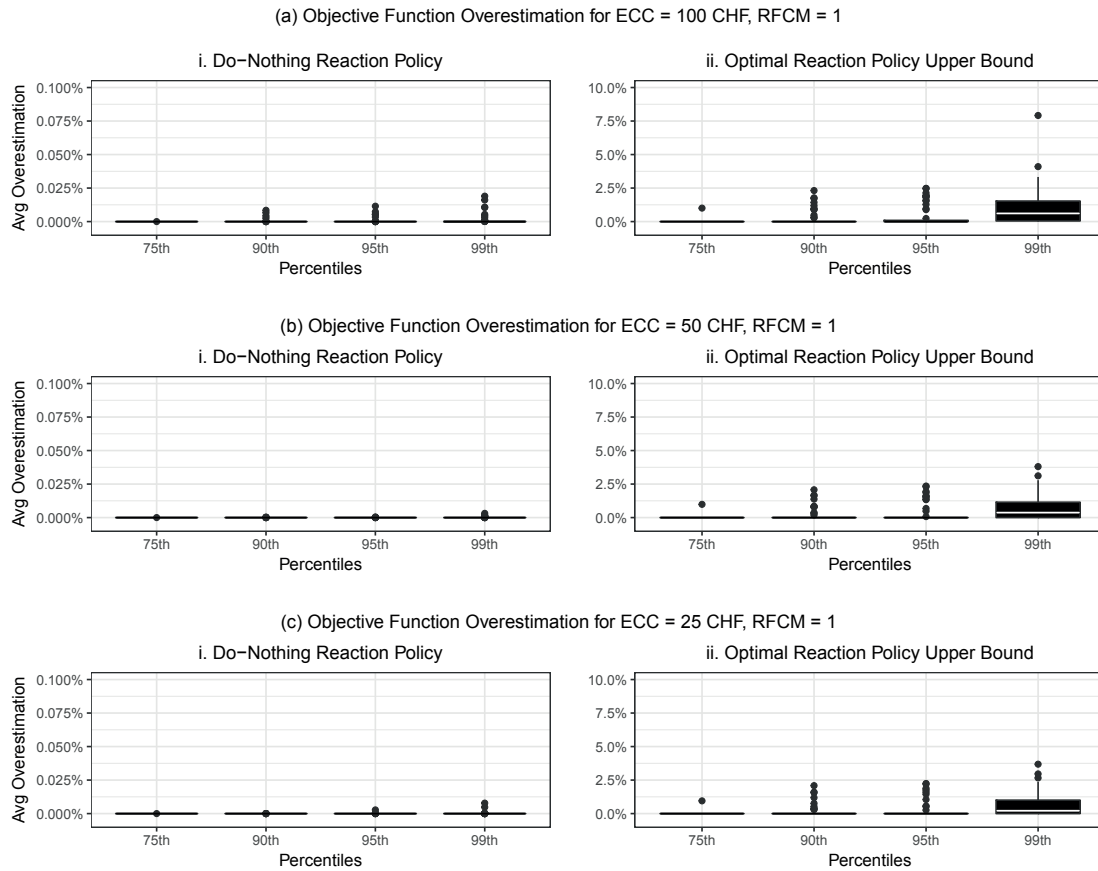
Assessing the Objective Function's Overestimation of the Real Cost.

In Section 4.4.1, we discussed the objective function's overestimation of the real cost, which is due to the un-captured effect in most parts of the objective function of demand points stocking out earlier than expected. To study this effect, we perform the simulation experiment described in Section 3.4.3 of Chapter 3 counting the number of realized overflows. Given the final solution of each instance, we simulate 10,000 scenarios, sampling the independent normally distributed error terms ε_{it} for each container $i \in \mathcal{P}$ and each day $t \in \mathcal{T}$, and applying them to the expected demand $\mathbb{E}(\rho_{it})$. Then, for each scenario, we analyze the effect of realized overflows on the objective function's overestimation of the real cost.

Computing the overestimation due to a do-nothing reaction policy is trivial. In the absence of inventory holding costs, which is the case for the waste collection IRP instances, the effect is only present in the ERFC component. Clearly, the overestimation will be higher for an optimal reaction policy which, in the occurrence of overflows, re-optimizes all subsequent decisions. However, computing the overestimation due to an optimal reaction policy has a significant computational burden, as it requires that re-optimization be done after each overflow for the 10,000 simulated scenarios. Therefore, we consider the following intuitive upper bound on the overestimation due to an optimal reaction policy. Consider a container i that overflows on day g and is visited for a regular collection on days $t > g$. Now, take the minimum day $h = \min t > g$ on which container i is visited for a regular collection and posit an optimal reaction policy so good that it removes the cost effect of container i from all days $t \geq h$. In other words, 1) we remove the container from all tours performed on days $t \geq h$. This implies the highest possible overestimation of the RC and ERFC components. 2) We also remove the probability of overflow on days $t \geq h$, which implies the highest possible overestimation of the ESOEDC component.

Figure 4.4 plots the overestimation for a do-nothing reaction policy as well as the discussed upper bound on the overestimation for an optimal reaction policy of objective (4.55) at the 75th, 90th, 95th and 99th percentile of the 10,000 scenarios. We present the results for an Emergency Collection Cost (ECC) $\zeta_i = 25, 50$ and 100 CHF, identical for all containers, and a Route Failure Cost Multiplier (RFCM) $\psi = 1$. Each box-plot is constructed using the average values over 10 runs for each of the 63 instances. The overestimation for the do-nothing reaction policy is marginal, which is due to the low probability of route failure observed in general for the waste collection IRP instances (see Section 3.4.3 in Chapter 3). Unsurprisingly, the upper bound on the overestimation for the optimal reaction policy appears to be linked to the level of the ECC. The median upper bound is approximately zero for the 75th, 90th and 95th percentile, with the maximum values reaching 2.5%. It becomes more pronounced at the 99th percentile, where the median values are 0.61%, 0.37% and 0.22% for an ECC of 100 CHF, 50 CHF and 25 CHF, respectively, which indicates the generally very low level of overestimation of the real cost. The maximum values do not exceed 8% for an ECC of 100 CHF, and 4% for an ECC of 50 CHF and 25 CHF. There is a strong correlation in the order of 70% between the number of realized overflows and the upper bound across the 63 instances. In

Figure 4.4: Objective Function's Overestimation of the Real Cost



Section 4.3, we argued the importance of tractability in terms of the probability calculations that enter the objective function. Using simpler probability expressions ignores some of the uncertainty propagation which, as proved in Proposition 4.2, leads to an overestimation of the real cost. Nevertheless, the results here indicate that this overestimation is marginal, and even a straightforward bound on the optimal reaction policy implies a median overestimation of the real cost of less than 1%.

4.6.3 Solving the Facility Maintenance Problem

The facility maintenance problem, as defined in Section 4.5.5, considers a set of facilities that have to be periodically inspected in order to limit the occurrence of breakdowns. Unlike the waste collection IRP, this problem does not consider demands. Thus, there is no deterministic equivalent to the stochastic problem. We start by comparing the two stochastic models proposed in Section 4.5.5. The models (FMP1) and (FMP2) treat uncertainty using a probabilistic objective function and probabilistic constraints, respectively. While both approaches use the same probability information, they do not use it in the same way. Specifically, the

Chapter 4. A Unified Framework for Rich Routing Problems with Stochastic Demands

probabilistic objective approach calculates the probability of incurring the emergency repair cost and lets the model determine the best balance between the routing and the expected emergency repair cost. The breakdown probabilities in the final solution thus depend on the value of the emergency repair cost itself. The probabilistic constraints approach controls the probability of breakdown in a rather artificial way. One usually knows what it costs to perform an emergency repair, while it is unclear what a reasonable value of the maximum allowable probability of breakdown γ^{DP} should be. At any rate, while these two approaches are different modeling-wise and from a conceptual stance, they are expected to be able to produce the same range of results. To verify this, we solve the model (FMP1) for a set of Emergency Repair Cost (ERC) ζ_i values, where ζ_i is identical for each facility $i \in \mathcal{P}$, and the model (FMP2) for a set of values for γ^{DP} .

The results are summarized in Tables 4.3 and 4.4, where each line is an averaged result over the 94 instances. In both tables, the first column identifies the modeling approach, the second one reports the value of the ERC and the third one the value of the maximum allowable breakdown probability γ^{DP} . In Table 4.3, the fourth column presents the average runtime in seconds, while the fifth and sixth columns report the average number of tours and facility visits, respectively. The rest of the columns report the best, average and worst results over 10 runs, and the percent gap between the average and best, and the worst and best results. Computation times are reasonable and, as expected, strongly correlated to the number of facility visits, and as a result to the cost. Not surprisingly, higher numbers of facility visits also correspond to higher numbers of tours. The gap between the average and the best solutions is in the order of 1-2%, and the gap between the worst and the best solutions is in the order of 2-3%, evidence of the stability of the ALNS.

Table 4.4 decomposes the solution cost into Routing Cost (RC) and Expected Emergency Repair Cost (EERC), whose averages are provided in the fourth and fifth columns, respectively. The last four columns are the result of a simulation experiment with 10,000 scenarios as

Table 4.3: Basic Results for Model (FMP1) vs. Model (FMP2)

Model	ERC	γ^{DP}	Runtime (s.)	Avg Num Tours	Avg Num Visits	Best Cost (CHF)	Avg Cost (CHF)	Worst Cost (CHF)	Gap Avg- Best (%)	Gap Worst- Best (%)
(FMP1)	1000.00	–	585.81	3.18	51.90	1810.57	1831.88	1857.70	1.18	2.60
(FMP1)	500.00	–	558.97	2.98	45.88	1594.88	1618.29	1641.24	1.47	2.91
(FMP1)	250.00	–	508.93	2.51	39.35	1404.90	1421.62	1443.89	1.19	2.78
(FMP1)	100.00	–	419.05	1.81	27.19	1125.71	1139.90	1158.42	1.26	2.91
(FMP1)	50.00	–	484.82	0.87	12.70	852.41	853.69	855.13	0.15	0.32
(FMP1)	25.00	–	478.33	0.84	2.75	556.32	556.32	556.32	0.00	0.00
(FMP2)	–	0.25	248.72	0.84	2.31	195.73	195.73	195.73	0.00	0.00
(FMP2)	–	0.20	319.59	0.99	6.81	304.19	304.27	304.52	0.03	0.11
(FMP2)	–	0.15	410.53	1.37	19.91	575.17	576.80	579.06	0.28	0.68
(FMP2)	–	0.10	500.31	1.99	29.02	836.40	841.00	845.57	0.55	1.10
(FMP2)	–	0.08	550.26	2.27	36.31	1003.84	1010.44	1016.76	0.66	1.29
(FMP2)	–	0.05	584.48	2.62	41.07	1144.97	1154.82	1166.69	0.86	1.90
(FMP2)	–	0.04	584.95	2.86	41.63	1201.01	1212.19	1226.82	0.93	2.15
(FMP2)	–	0.03	627.74	2.88	43.83	1237.16	1249.44	1264.50	0.99	2.21
(FMP2)	–	0.02	667.10	3.49	49.01	1438.88	1453.06	1463.91	0.99	1.74

Table 4.4: Performance Indicators for Model (FMP1) vs. Model (FMP2)

Model	ERC	γ^{DP}	Avg RC (CHF)	Avg EERC (CHF)	Avg Num Breakdowns			
					75th Perc.	90th Perc.	95th Perc.	99th Perc.
(FMP1)	1000.00	–	1444.59	387.29	1.21	1.98	2.47	3.44
(FMP1)	500.00	–	1304.26	314.03	1.76	2.53	3.13	4.20
(FMP1)	250.00	–	1108.69	312.94	2.59	3.49	4.13	5.34
(FMP1)	100.00	–	780.78	359.12	5.20	6.55	7.41	9.07
(FMP1)	50.00	–	369.76	483.93	11.55	13.54	14.74	17.06
(FMP1)	25.00	–	201.93	354.39	16.46	18.76	20.18	22.84
(FMP2)	–	0.25	195.73	0.00	16.75	19.02	20.48	23.17
(FMP2)	–	0.20	304.27	0.00	14.35	16.50	17.82	20.36
(FMP2)	–	0.15	576.80	0.00	9.19	10.97	12.07	14.18
(FMP2)	–	0.10	841.00	0.00	5.62	6.98	7.85	9.58
(FMP2)	–	0.08	1010.44	0.00	3.91	5.06	5.84	7.29
(FMP2)	–	0.05	1154.82	0.00	2.53	3.48	4.11	5.31
(FMP2)	–	0.04	1212.19	0.00	2.17	3.06	3.58	4.75
(FMP2)	–	0.03	1249.44	0.00	2.01	2.82	3.41	4.52
(FMP2)	–	0.02	1453.06	0.00	1.22	2.02	2.47	3.43

the one previously described, and report the average number of breakdowns over the 94 instances at the 75th, 90th, 95th and 99th percentile of the 10,000 scenarios. There appears to be, as expected, a clear negative correlation between the routing cost and the number of breakdowns at any percentile. This happens because higher routing costs are associated with more frequent facility visits and, as per formula (4.64), with lower breakdown probabilities. Moreover, we notice that the routing cost and the number of breakdowns for models (FMP1) and (FMP2) vary within similar ranges. This is confirmed by Figure 4.5, which is a visual representation of the above results. It demonstrates that the two approaches are logically equivalent, with similar routing costs corresponding to similar levels of occurrence of breakdowns. We stress again that both approaches are probabilistic, using the same uncertainty information in different ways.

To complete the picture, we compare the two probabilistic models to model (FMPD) of Section 4.5.5, which is a flexible deterministic approach oblivious to any uncertainty information. It considers a routing-only objective function and stipulates that each facility $i \in \mathcal{P}$ must be visited at least ν_i times during the planning horizon. Table 4.5, which is structured in the same way as Table 4.4, summarizes the results for $\nu_i = 1$ and 2, with ν_i identical for all $i \in \mathcal{P}$. Some of the instances become infeasible for higher values of ν_i . The table clearly shows that it takes a much higher routing cost to achieve similar levels of occurrence of breakdowns, thus highlighting the superiority of the stochastic modeling approaches. Similar to Section 3.4.3 of Chapter 3, we compare the performance of two approaches in terms of number of breakdowns at different percentiles. Thus, we isolate the EERC component from the solution cost of the probabilistic models, and compare only the RC components.

Chapter 4. A Unified Framework for Rich Routing Problems with Stochastic Demands

Figure 4.5: Comparison of Routing Cost and Breakdowns for Model (FMP1) vs. Model (FMP2)




Table 4.5: Performance Indicators for Model (FMPD)

Model	ERC	ν_i	Avg RC (CHF)	Avg EERC (CHF)	Avg Num Breakdowns			
					75th Perc.	90th Perc.	95th Perc.	99th Perc.
(FMPD)	-	2	1945.96	0.00	3.16	4.10	4.56	5.71
(FMPD)	-	1	1140.10	0.00	4.28	5.47	6.26	7.77

4.7 Summary

This chapter introduces, analyzes and formulates a unified framework for modeling and solving various classes of rich routing problems. Demand is stochastic, can be non-stationary, and is forecast with any model that provides the expected demands over the planning horizon and the distribution of the error terms. The formulation includes many rich routing features relevant to real-world problems, such as multiple depots, open and multi-period tours, intermediate facilities, time windows, maximum tour duration, accessibility restrictions, visit periodicities and service choice, workload balancing, continuity of service, etc.

The practical applicability of the approach is reinforced by the fact that the probability information related to demand stochasticity can be pre-computed or at least partially pre-processed. We highlight the fact that the stock-out/overflow/breakdown probabilities can be pre-computed for error terms from any distribution and with any covariance structure among them. The assumption of iid error terms, still from any distribution, remains necessary for partially pre-processing the route failure probabilities. The last restriction allows us to

preserve tractability, which is critical for operational problems such as those discussed in the text. Finally, we show that certain problems where the inventory component is not present, such as facility maintenance, can still be viewed through the prism of inventory routing, with event probabilities at the demand points, or breakdown probabilities in this specific example, accumulating as would inventory.

Using the waste collection IRP instances from Chapter 3, we demonstrate that pre-processing the route failure probabilities through the derivation of ECDFs under the assumption of iid error terms is sufficient to preserve tractability. Simulating the error terms on the final solution further allows us to verify the low level of occurrence of overflows and shows that the objective is an excellent representation of the real cost. On the new set of facility maintenance instances, our framework is able to achieve the same level of occurrence of breakdowns for a significantly lower routing cost compared to alternative deterministic policies.

5 Conclusion

This chapter borrows from the articles:

Markov, I., Varone, S., and Bierlaire, M. (2016). Integrating a heterogeneous fixed fleet and a flexible assignment of destination depots in the waste collection VRP with intermediate facilities, *Transportation Research Part B: Methodological* 84:256-273.

Markov, I., Bierlaire, M., Cordeau, J.-F., Maknoon, Y., and Varone, S. (2016). Inventory routing with non-stationary stochastic demands. Technical report TRANSP-OR 160825. Transport and Mobility Laboratory, EPFL, Lausanne, Switzerland.

Markov, I., Bierlaire, M., Cordeau, J.-F., Maknoon, Y., and Varone, S. (2017). A general framework for routing problems with stochastic demands. *Proceedings of the 17th Swiss Transport Research Conference (STRC)*, May, 17-19, 2017.

The work therein has been performed by the author in collaboration with Prof. Michel Bierlaire, Prof. Jean-François Cordeau, Prof. Yousef Maknoon and Prof. Sacha Varone.

In this chapter, Section 5.1 revisits the objectives and contributions, and analyzes how they are supported by the main findings in Chapters 2, 3 and 4. Given the direct applicability of our work, Section 5.2 examines its practical implications and the challenges we may face in the way of its widespread adoption. Finally, Section 5.3 closes by identifying and discussing promising areas of future research.

5.1 Main Findings

In this thesis, we propose a unified framework for modeling and solving various classes of rich routing problems with stochastic demands, including among others the VRP and the IRP. We solve numerous challenges related to the realistic modeling of demand uncertainty,

its effect on undesirable events and the costs of their associated recourse actions. These elements are used in the development of a tractable approach and a powerful solution methodology, successfully applied on a real case study inspired by the problem of collecting recyclables from sensorized containers in the canton of Geneva, Switzerland. Chapter 2 starts with the deterministic single-day problem, a rich VRP with intermediate facilities. For it, we build an MILP formulation, which is enhanced with a number of valid inequalities and includes a variety of rich routing features, in particular a heterogeneous fixed fleet and the possibility of open tours. Moreover, it considers a general cost function corresponding to the cost structure of a typical firm. To solve realistic instances, we develop a meta-heuristic approach based on multiple neighborhood search.

The extensive computational testing confirms the benefit of our valid inequalities to the optimization model. The meta-heuristic approach achieves optimality on small instances, exhibits competitive performance in comparison to state-of-the-art solution methods for special cases of our problem, and leads to important savings in the state of practice. It presents fast computation times and outperforms significantly the solution currently in place at the collector in our case study in terms of quality and functionality. We are also able to show that the possibility of open tours with different origin and destination depots can lead to noticeable savings, especially in rural and sparsely populated areas where such benefits will be most pronounced.

In Chapter 3, we extend the single-day problem to a finite horizon, which results in the waste collection IRP. Here, demand is stochastic, can be non-stationary, and is forecast using any model that provides the expected demands over the planning horizon and a measure of uncertainty represented by the standard deviation of the error terms, the latter assumed to be iid normal. The objective captures demand uncertainty with the goal of minimizing the expected cost, including the expected cost of recourse actions, subject to a range of practical and policy-related constraints. To manage the increased complexity, we build a powerful Adaptive Large Neighborhood Search (ALNS) algorithm which produces excellent results on VRP and IRP benchmarks sets from the literature.

The computational experiments demonstrate that including probabilistic information in the objective function leads to a relatively modest increase in the routing cost, while avoiding major expenditures that otherwise occur even at moderate percentiles of the simulated demand realization scenarios. The probabilistic approach significantly outperforms alternative deterministic policies of using artificially low capacities for the containers and the trucks in its ability to limit the occurrence of container overflows for the same routing cost. We also analyze the solution properties of a rolling horizon approach in terms of empirical lower and upper bounds. Our results show the benefit of a dynamic stochastic approach, which includes new information at each rollover, in comparison to a static stochastic approach over the same planning horizon.

Chapter 4 generalizes the approach in a unified framework for rich routing problems with sto-

chastic demands. The formulation includes many rich routing features relevant to real-world problems, such as multiple depots, open and multi-period tours, intermediate facilities, time windows, maximum tour duration, accessibility restrictions, visit periodicities and service choice, workload balancing, continuity of service, etc. Demand is stochastic, can be non-stationary, and is now forecast using any model that provides the expected demands over the planning horizon and the distribution of the error terms. We relax fully or partially the assumption of iid normal error terms, consider a general inventory policy, discuss tractability related topics, and illustrate applications to a variety of rich routing problems borrowed from the literature and inspired from practice. In particular, we show that certain problems where the inventory component is not present, such as facility maintenance, can still be viewed through the prism of inventory routing, with event probabilities at the demand points, or breakdown probabilities in this specific example, accumulating as would inventory.

The computational experiments focus on the topic of complexity vs. tractability, indicating that the modeling simplifications we use for preserving tractability do not compromise our representation of the real cost. The numerical experiments on a new set of facility maintenance instances confirm the conclusions from the waste collection IRP instances of the superiority of the stochastic approach in comparison to alternative deterministic policies.

5.2 Practical Implications

Waste collection is one of the most important logistical activities performed by any municipality, and also one of the most expensive. According to various estimates, collection costs account for more than 60% of waste management costs (Johansson, 2006; Tavares et al., 2009; Greco et al., 2015; Asimakopoulos et al., 2016). Recycling, on the other hand, can alleviate problems related to landfill capacity and pollution, and many countries have already set ambitious target levels for recycling. As part of its Circular Economy Strategy, the European Union, for example, has adopted legislative proposals to set a common target for recycling 65% of municipal and 75% of packaging waste by 2030, limiting at the same time the use of landfills (European Commission, 2016). Given the high cost of waste management and the significant proportion of collection costs, even small improvements in the latter can lead to substantial financial savings for waste collectors, municipalities, and ultimately the taxpayer.

In this context, waste management firms are often resistant to change, with collection traditionally based on fixed tours executed with a daily or weekly frequency and little to no regard for optimality. That being said, the market is becoming more competitive. In the Geneva area alone, there are multiple collectors of recyclable materials. They invoice the municipality based on collected volume. Thus, the efficient fleet utilization and routing, and the intelligent scheduling of container visits are of paramount importance. Our framework can easily be embedded in a Geographic Information System (GIS), with an interactive user

interface for problem definition and solution display. Moreover, it is straightforward to include some user input in the solution, such as intentionally skipping or visiting certain containers. The GIS would allow plotting the solution with shortest or fastest paths among containers and, most importantly, exporting it in a GPS readable format on a tablet or another device that the truck driver uses to navigate. Unfortunately, such high-tech collectors are still not the norm in most parts of the world, but the application of intelligent waste collection is gaining traction, evidenced by the number of pilot projects involving waste collection GIS and sensorized containers (see e.g. Ghose et al., 2006; Oliveira Simonetto and Borenstein, 2007; Krikke et al., 2008; Repoussis et al., 2009; Rovetta et al., 2009; Tavares et al., 2009; Zamorano et al., 2009; Arribas et al., 2010; Faccio et al., 2011; McLeod et al., 2013; Mes, 2012; Anghinolfi et al., 2013; Mes et al., 2014; Christodoulou et al., 2016). Beliën et al. (2014) provide a comprehensive review of municipal solid waste collection and management problems.

An interesting practical observation from our work with actors in the waste management industry is the importance of the ability to explain what the models and algorithms do. Intuitive and commonsense approaches have better chances of being adopted in practice. Our framework is based on a sophisticated modeling and solution approach, yet the concepts behind it—probabilities, costs, undesirable events—and the relationships among them defined in the objective and constraints are universal and easy to understand. From a more general point of view, our unified framework is an effort to bridge the gap between theory and practice, highlighted and discussed by Gendreau et al. (2016) in their survey of stochastic vehicle routing problems. We include rich practice-driven routing features and relax commonly used assumptions in the literature, such as iid normal random variables or demand stationarity. Finally, we integrate real-world demand forecasting models and show that the resulting framework becomes operational and provides excellent results on real case studies.

5.3 Future Research Directions

This thesis treats a complex real-world problem, integrating techniques from optimization, statistics and simulation, and relying on efficient algorithmic implementations. Thus, it lends itself to a wide array of potential future work directions. We can classify them into two groups: those of mainly practical interest and those of mainly theoretical interest. Starting from the first group, in our view the most important task is the development of additional benchmark instances, which will allow us to test the framework's full capabilities on different problem types. While there exist benchmark instances for many of the reviewed problems, in particular those modeled in Chapter 4, they are largely deterministic or involve simplistic routing structures. Thus, it is crucial that the instances be based on, or at least derived from, real data. We are interested in evaluating how the framework performs on concrete problems faced by real actors in the transportation field.

The optimization models presented in this work already include a variety of rich routing features. Still, other practically relevant ideas of particular interest may be the generaliza-

tion of single to multiple time windows or the integrated solution of the multi-commodity problem. A further extension of the latter is compartmentalization, which allows vehicles to deliver or collect more than one commodity at a time (Mendoza et al., 2011). In Switzerland, for example, glass is often collected separately as white and colored glass. Since both types are usually recycled at the same facility, they can be collected by the same vehicle. Flexible compartmentalization where the compartment separators can be adjusted after each vehicle emptying offer the possibility of further savings. Moreover, in the absence of specific requirements, any commodity can be transported by any vehicle. Since a vehicle may perform multiple intermediate facility visits during a tour, it may transport a different commodity after each visit. Finally, this can be considered in a competitive market where suppliers or collectors compete for customers, and where undesirable events such as stock-outs or overflows may lead to customers switching to the competition. Different strategic behaviors, including collaboration, may be applicable. These extensions would undoubtedly lead to complications in the solution methodology, with the need for developing new and more sophisticated operators for the ALNS, or any other methodology deemed appropriate.

In our framework, demand is modeled at discrete time periods and inventory is updated at the start of each period, which is not necessarily the case in reality. Indeed, visiting a customer at different times on a given day will probably imply different delivery quantities. Furthermore, the absence of a stock-out at the start of the day does not prevent the possibility of stock-out later that same day. Real-time, or online, optimization can be used to deal with continuous time demands. In this setup, the system is updated after each customer visit and the subsequent decisions are fully or partially re-optimized based on the latest available information, given the observed inventories of the already visited customers. Alternatively, the availability of more frequent sensor information can be exploited using a rolling horizon approach over a finer time discretization, for example an hour. This leads us to the second group of future work directions—those of primarily theoretical interest. From a more general perspective, our framework opens the door to developing even more comprehensive objective functions, capturing further probability propagations. This should make it possible to relax some of the assumptions and allow for increasingly complex routing structures. Finally, integration of other stochastic elements, in particular travel and service times, is extremely relevant, especially in application areas such as maritime routing. Finding the balance between modeling realism and the preservation of tractability is one of the main challenges in this direction.

In terms of solution methodology, rich routing problems of realistic size preclude the use of fully exact approaches. At the same time, the performance of meta-heuristics can only be evaluated on special cases of these problems, for which benchmark instances with known optimal solutions exist. Thus, an important future work area concerns the development of theoretical lower bounds. Our optimization models rely on arc-based formulations which are known to provide weak lower bounds (Semet et al., 2014; Poggi and Uchoa, 2014). Such models quickly become intractable even for moderate instance sizes. A promising direction is the development of a path-based formulation and a state-of-the-art column generation

Chapter 5. Conclusion

procedure. While this is rather straightforward for the linear and deterministic case where the pricing problem is an Elementary Shortest Path Problem with Resource Constraints (ESPPRC), the non-linear nature of our objective function will certainly pose challenges in this regard. Certain simplifications in the routing structure and approximations like demand discretization may be sufficient to linearize the objective function and cast the pricing problem as an ESPPRC. Alternatively, a more complicated pricing problem will need to be modeled and solved. At any rate, even if the column generation approach is unable to solve realistic-size instances, it may still be capable of providing good-quality lower bounds on the ALNS, in addition to the encouraging indications we already have in terms of its stability in general and its performance on classical VRP and IRP benchmarks from the literature.

A Waste Collection IRP: Additional Analysis

A.1 Waste Collection IRP: Collection Strategies

Tables A.1, A.2 and A.3 below report the average level, the average level at collection, and the number of instances with container collections, by day, respectively. Each row in the tables is an averaged result over 10 runs of the 63 instances for complete solutions with different Emergency Collection Costs (ECC) and Route Failure Cost Multipliers (RFCM), and for the routing-only solution. In the tables, the first three columns identify the objective, i.e. complete vs. routing-only, the ECC and the RFCM, while the rest of the columns report the relevant statistics for each table for $t \in \mathcal{T} \cup \mathcal{T}^+$.

Starting from Table A.1, we note that the average level on day $t = 0$ is independent of any action and is thus the same for each combination of ECC and RFCM. For the rest of the days, the principal difference is between the series of complete solutions on the one hand and the routing-only solution on the other. The average container level of the routing-only solution is often twice as high as that of the complete solutions. With the exception of day $t = 0$, the average level of containers in the complete solution is in the order of 20 or 30%, while that

Table A.1: Average Level of All Containers by Day

Objective	ECC	RFCM	Avg level $t = 0$ (%)	Avg level $t = 1$ (%)	Avg level $t = 2$ (%)	Avg level $t = 3$ (%)	Avg level $t = 4$ (%)	Avg level $t = 5$ (%)	Avg level $t = 6$ (%)	Avg level $t = 7$ (%)
Complete	100.00	1.00	56.40	23.65	27.51	32.16	27.38	24.65	30.71	35.93
Complete	100.00	0.50	56.40	23.57	27.19	32.08	27.66	24.54	30.61	35.82
Complete	100.00	0.25	56.40	23.53	27.03	32.09	28.15	24.49	30.55	35.76
Complete	100.00	0.00	56.40	23.49	27.24	32.03	27.82	24.55	30.61	35.82
Complete	50.00	1.00	56.40	25.45	27.28	32.50	25.72	26.58	32.64	37.85
Complete	50.00	0.50	56.40	25.68	27.43	32.65	25.11	26.58	32.64	37.85
Complete	50.00	0.25	56.40	25.74	27.41	32.66	25.43	26.51	32.57	37.79
Complete	50.00	0.00	56.40	25.62	27.39	32.62	25.44	26.50	32.56	37.77
Complete	25.00	1.00	56.40	26.38	27.42	32.60	25.14	27.83	33.89	39.10
Complete	25.00	0.50	56.40	26.41	27.42	32.49	24.97	27.80	33.87	39.08
Complete	25.00	0.25	56.40	26.50	27.50	32.64	24.99	27.81	33.87	39.08
Complete	25.00	0.00	56.40	26.49	27.45	32.50	24.79	27.79	33.85	39.06
Routing-only	0.00	0.00	56.40	53.40	41.82	42.82	48.01	53.90	59.96	65.17

Appendix A. Waste Collection IRP: Additional Analysis

of the routing-only solution is in the order of 40, 50 or 60%, undoubtedly contributing to the much higher number of container overflows reported in Table 3.8. As far as the series of complete solutions is concerned, we observe some minor differences suggesting in most cases that the average level of containers is inversely correlated to the ECC.

In Table A.2, we see that the differences in the level of collected containers are less pronounced for the first two days and become progressively higher later in the planning horizon. On days $t = 3$ and $t = 4$, the complete solutions collect containers that are on average less than half full in order to minimize the probability of overflows on days $t = 5$, $t = 6$ and $t = 7$, where there are no collections. On the other hand, the routing-only solution collects containers that are on average 70 to 80% full, focusing only on those containers that are expected to overflow in the planning horizon, and ignoring all the rest.

Table A.3 reports for each day of the planning horizon the number of instances, from a total of 63, which collect containers on that day. While we do not see much difference for day $t = 0$, the differences for the rest of the planning horizon are consequential. The complete

Table A.2: Average Level of Collected Containers by Day

Objective	ECC	RFCM	Avg level $t = 0$ (%)	Avg level $t = 1$ (%)	Avg level $t = 2$ (%)	Avg level $t = 3$ (%)	Avg level $t = 4$ (%)	Avg level $t = 5$ (%)	Avg level $t = 6$ (%)	Avg level $t = 7$ (%)
Complete	100.00	1.00	64.12	61.81	63.31	44.27	41.63	-	-	-
Complete	100.00	0.50	64.23	61.56	63.82	44.37	42.21	-	-	-
Complete	100.00	0.25	64.10	61.53	63.26	44.39	42.42	-	-	-
Complete	100.00	0.00	64.09	61.51	64.71	44.15	42.42	-	-	-
Complete	50.00	1.00	64.73	61.85	57.95	44.20	41.90	-	-	-
Complete	50.00	0.50	64.80	61.74	59.33	43.90	42.54	-	-	-
Complete	50.00	0.25	64.76	61.69	-	44.10	42.81	-	-	-
Complete	50.00	0.00	64.78	61.65	59.43	44.03	42.47	-	-	-
Complete	25.00	1.00	65.19	61.84	55.70	44.39	41.82	-	-	-
Complete	25.00	0.50	65.19	61.84	57.33	44.27	40.26	-	-	-
Complete	25.00	0.25	65.26	61.72	53.45	44.38	42.07	-	-	-
Complete	25.00	0.00	65.25	61.72	55.92	44.09	41.99	-	-	-
Routing-only	0.00	0.00	78.73	72.68	78.32	78.76	71.60	-	-	-

Table A.3: Number of Instances with Container Collections by Day

Objective	ECC	RFCM	Num for $t = 0$	Num for $t = 1$	Num for $t = 2$	Num for $t = 3$	Num for $t = 4$	Num for $t = 5$	Num for $t = 6$	Num for $t = 7$
Complete	100.00	1.00	60	3	2	47	45	-	-	-
Complete	100.00	0.50	60	3	2	44	49	-	-	-
Complete	100.00	0.25	60	3	2	45	49	-	-	-
Complete	100.00	0.00	60	3	2	46	47	-	-	-
Complete	50.00	1.00	59	6	1	50	31	-	-	-
Complete	50.00	0.50	59	6	1	52	33	-	-	-
Complete	50.00	0.25	59	6	0	52	34	-	-	-
Complete	50.00	0.00	59	6	1	53	33	-	-	-
Complete	25.00	1.00	57	6	2	49	21	-	-	-
Complete	25.00	0.50	57	6	3	51	19	-	-	-
Complete	25.00	0.25	57	6	3	50	24	-	-	-
Complete	25.00	0.00	57	6	4	53	21	-	-	-
Routing-only	0.00	0.00	52	60	36	6	4	-	-	-

A.2. Waste Collection IRP: Effect of Lower Truck Capacity on the Solution Cost

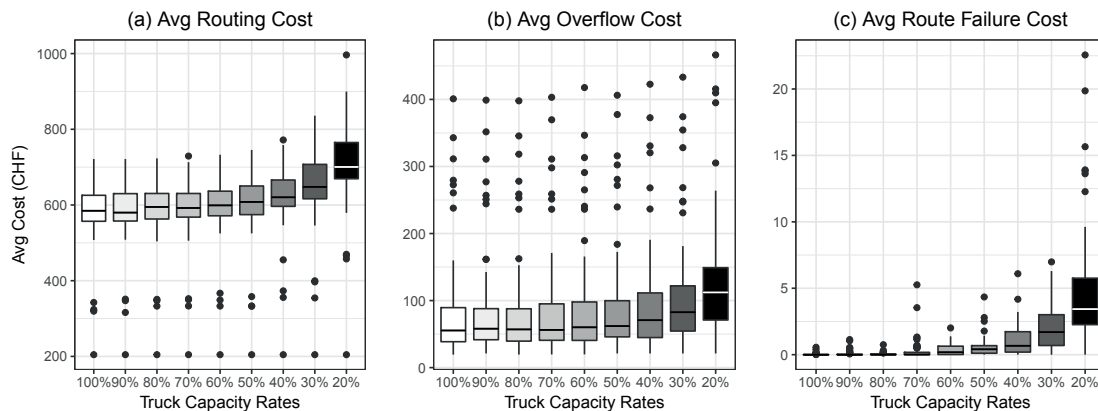
solutions seem to postpone collections until days $t = 3$ and $t = 4$ in order to minimize the probability of overflow on days $t = 5$, $t = 6$ and $t = 7$, where the vehicles are not available for performing collections. The routing-only solution, being completely oblivious to the probability of overflow, performs most of the collections until day $t = 2$, thus allowing containers to be fuller later in the planning horizon. This is a clear example of the lack of foresight in the routing-only solution, and why it results in a much higher number of overflows. As far as the series of complete solutions is concerned, the solutions with a higher ECC seem to perform more collections and shift collections later in the planning horizon.

A.2 Waste Collection IRP: Effect of Lower Truck Capacity on the Solution Cost

The results in Tables 3.7 and 3.8, and Figure 3.6 indicate the low level of occurrence overflows and the associated marginal contribution of the ERFC component to the total solution cost. We perform additional analysis to study this effect by rerunning the experiments on the 63 waste collection IRP instances using lower capacity trucks. For each instance, we reduce the truck capacities to $n\%$ of their original values, where n varies from 100% to 20% by a step of 10%. Some of the instances become infeasible for $n < 20\%$. We solve the problem with the complete objective function for an ECC of 100 CHF and an RFCM of one. The goal is to analyze how the objective function and its components, in particular the ERFC component, react to lowering the truck capacity.

Parts (a), (b) and (c) of Figure A.1 below plot the effect of lowering the truck capacities on the average routing cost, the average overflow cost, and the average route failure cost, respectively. We observe that the effect is most noticeable for $n < 50\%$, suggesting that the original vehicle capacities are quite high relative to the container demands. Noting the differences in the scales of the y-axes, we observe that the highest nominal increase is in the average routing cost, while the highest relative increase is in the average route failure

Figure A.1: Effect of Lower Truck Capacity on the Solution Cost




cost. The latter goes from practically zero to a median value of 3.43 CHF, with the maximum reaching 22.56 CHF. While this increase is important in relative terms, the ERFC component still remains a small fraction of the total cost. This is due to the fact that tours adapt to lower vehicle capacities by making more dump visits on average, thus partially shifting of the effect to the routing cost.

A.3 Waste Collection IRP: Effect of Open Tours on the Solution Cost

For the waste collection VRP in a deterministic setting, Chapter 2 finds an average improvement of 2.54% when allowing open tours with an optimization of the home depot on modifications of the Crevier et al. (2007) instances. The improvement exceeds 10% for some of the instances, and appears to be negatively correlated with the instance size. To evaluate this effect in a multi-day setting, we add additional depots to the 63 waste collection IRP instances. The number of added depots is specific to each vehicle and relies on the case study data. In the resulting instances, each vehicle is allowed to visit up to 3 depots, one of them being its home depot. All vehicles are free to end their tours at any depot on day t and start from this depot on day $t + 1$, and are only required to return to their home depots on Friday. We analyze the cost benefit of open tours on the complete objective with an ECC of 100 CHF and an RFCM of one, and on the routing-only objective.

Figure A.2 plots the best results over 10 runs for each instance, where part (a) depicts the effect on the routing cost and part (b) on the overflow cost. The contribution of the route failure cost component to the total cost is immaterial and is ignored in the analysis. The results indicate the clear cost benefit of open tours. For both the complete and the routing-only objective function, allowing open tours leads to an average decrease of the routing cost of approximately 10%. The effect holds virtually across all instances as visible in part (a) of Figure A.2. It is less pronounced in the case of the overflow cost for the complete objective function. Nevertheless, while less noticeable in absolute terms, the average relative decrease is again approximately 10%. The case study in Chapter 2 mentions regions where open tours

Figure A.2: Effect of Open Tours on the Solution Cost



A.3. Waste Collection IRP: Effect of Open Tours on the Solution Cost

with destination depots different from the origin depots are practiced. Therefore, this result demonstrates that the findings and conclusions therein are valid, and in fact even stronger, for a multi-day problem. The improvements for a multi-day problem do not seem to be related to the instance size.

The experiments in this section use the ALNS of Section 3.3 of Chapter 3. To handle multiple depots, we add the following operator to the list of repair operators in Section 3.3.2:

- *Replace a destination depot:* This operator selects a random tour and replaces its destination depot with a random destination depot $o \in \mathcal{O}''_{kt}$, where $t \in \mathcal{T}$ is the period in which the tour is executed and $k \in \mathcal{K}$ is the vehicle executing it. The algorithm then finds $h = \min t' > t$ s.t. $H_{kh} > 0$, i.e. the next period h for which vehicle k is available, and changes the origin depot of the tour that vehicle k executes in period h to o .

B Equivalence of Stock-out and Overflow Probabilities

At the demand point level, the undesirable event for a distribution problem is a stock-out, while for a collection problem it is an overflow. Here, we prove the following:

Proposition B.1. *The calculation of the probability of overflow for a collection problem is identical to the calculation of the probability of stock-out for a distribution problem.*

Proof. Let Λ'_{ig} denote the inventory after a regular collection of demand point i in period g . This collection is accompanied by a corresponding delivery of empty space. Thus, the empty space inventory after a regular delivery is $\Lambda_{ig} = (\omega_i - \Lambda'_{ig})$, where ω_i is the capacity of demand point i . Given a regular collection in period g , the unconditional probability of overflow of demand point i in period $g + 1$ is expressed as:

$$\mathbb{P}(\Lambda'_{ig} + \rho_{ig} \geq \omega_i) = \mathbb{P}((\omega_i - \Lambda'_{ig}) - \rho_{ig} \leq 0) = \mathbb{P}(\Lambda_{ig} - \rho_{ig} \leq 0), \quad (\text{B.1})$$

the last expression being equivalent to expression (4.4) for a distribution problem. Given a regular collection in period g , the conditional probability of overflow in periods $d > g + 1$ is expressed as:

$$\begin{aligned} & \mathbb{P}\left(\Lambda'_{ig} + \sum_{t=g}^h \rho_{it} \geq \omega_i \mid \Lambda'_{ig} + \sum_{t=g}^{h-1} \rho_{it} < \omega_i\right) = \\ & \mathbb{P}\left((\omega_i - \Lambda'_{ig}) - \sum_{t=g}^h \rho_{it} \leq 0 \mid (\omega_i - \Lambda'_{ig}) - \sum_{t=g}^{h-1} \rho_{it} > 0\right) = \\ & \mathbb{P}\left(\Lambda_{ig} - \sum_{t=g}^h \rho_{it} \leq 0 \mid \Lambda_{ig} - \sum_{t=g}^{h-1} \rho_{it} > 0\right), \quad \forall h > g, \end{aligned} \quad (\text{B.2})$$

the last expression being equivalent to expression (4.5) for a distribution problem. The proofs for the unconditional and conditional probabilities of overflow given an emergency collection in period $g' > g$ follow as special cases. \square

C Processing Stochastic Information

C.1 Pre-computing the Stock-out Probabilities

To pre-compute the unconditional and conditional probabilities of stock-out (4.4)–(4.7), choose a sufficiently large number M and for $m \in \{1, \dots, M\}$ simulate:

$$\mathbf{e}^m = (e_{11}^m, \dots, e_{1|\mathcal{T}|}^m, e_{21}^m, \dots, e_{|\mathcal{P}||\mathcal{T}|}^m), \quad (\text{C.1})$$

by drawing $\boldsymbol{\varepsilon}$ from Φ , where $\boldsymbol{\varepsilon}$ is the vector of error terms defined by equation (4.2). Using the result of (C.1), the probability in formula (4.4) is pre-computed as:

$$\mathbb{P}(\Lambda_{ig} - \rho_{ig} \leq 0) = \mathbb{P}(\boldsymbol{\varepsilon}_{ig} \geq \Lambda_{ig} - \mathbb{E}(\rho_{ig})) = \frac{\sum_{m=1}^M \text{IF}(e_{it}^m \geq \Lambda_{ig} - \mathbb{E}(\rho_{ig}), 1, 0)}{M}. \quad (\text{C.2})$$

Using the same technique, the probability in formula (4.5) develops and pre-computes as:

$$\begin{aligned} & \mathbb{P}\left(\Lambda_{ig} - \sum_{t=g}^h \rho_{it} \leq 0 \mid \Lambda_{ig} - \sum_{t=g}^{h-1} \rho_{it} > 0\right) = \\ & = \mathbb{P}\left(\sum_{t=g}^h \boldsymbol{\varepsilon}_{it} \geq \Lambda_{ig} - \sum_{t=g}^h \mathbb{E}(\rho_{it}) \mid \sum_{t=g}^{h-1} \boldsymbol{\varepsilon}_{it} < \Lambda_{ig} - \sum_{t=g}^{h-1} \mathbb{E}(\rho_{it})\right) = \\ & = \frac{\mathbb{P}\left(\sum_{t=g}^h \boldsymbol{\varepsilon}_{it} \geq \Lambda_{ig} - \sum_{t=g}^h \mathbb{E}(\rho_{it}), \sum_{t=g}^{h-1} \boldsymbol{\varepsilon}_{it} < \Lambda_{ig} - \sum_{t=g}^{h-1} \mathbb{E}(\rho_{it})\right)}{\mathbb{P}\left(\sum_{t=g}^{h-1} \boldsymbol{\varepsilon}_{it} < \Lambda_{ig} - \sum_{t=g}^{h-1} \mathbb{E}(\rho_{it})\right)} = \\ & = \frac{\sum_{m=1}^M \text{IF}\left(\sum_{t=g}^h e_{it}^m \geq \Lambda_{ig} - \sum_{t=g}^h \mathbb{E}(\rho_{it}) \text{ AND } \sum_{t=g}^{h-1} e_{it}^m < \Lambda_{ig} - \sum_{t=g}^{h-1} \mathbb{E}(\rho_{it}), 1, 0\right)}{\sum_{m=1}^M \text{IF}\left(\sum_{t=g}^{h-1} e_{it}^m < \Lambda_{ig} - \sum_{t=g}^{h-1} \mathbb{E}(\rho_{it}), 1, 0\right)}, \quad (\text{C.3}) \\ & \quad \forall h > g. \end{aligned}$$

The function $\text{IF}([\text{condition}], 1, 0)$ is equal to 1 if the condition is satisfied, and to 0 otherwise. The probabilities in formulas (4.6) and (4.7) are pre-computed as special cases of (C.2) and (C.3), respectively. The time complexity of calculating each probability is linear in M .

C.2 Partially Pre-processing the Route Failure Probabilities

Formula (4.10), which defines the probability of route failure, develops as:

$$\begin{aligned}
 & \mathbb{P}(\Gamma_{\mathcal{S}} > \Omega_k) = \\
 & = \mathbb{P}(\mathbb{E}(\Gamma_{\mathcal{S}}) + \mathcal{E} > \Omega_k) = \\
 & = \mathbb{P}(\mathcal{E} > \Omega_k - \mathbb{E}(\Gamma_{\mathcal{S}})),
 \end{aligned} \tag{C.4}$$

where the cumulative error term \mathcal{E} is derived from the definition of the delivery quantity $\Gamma_{\mathcal{S}}$ in trip \mathcal{S} in formula (4.9) as follows:

$$\mathcal{E} = \sum_{t \in \mathcal{T} \setminus 0} \sum_{\mathcal{S}_t \in \mathcal{S}} \sum_{s \in \mathcal{S}_t} \sum_{h=m}^{t-1} \varepsilon_{sh}. \tag{C.5}$$

In the general case, the distribution of \mathcal{E} is unknown. And while probabilities (C.4) can be approximated using the simulation techniques presented in Appendix C.1, the number of combinations involving different periods, demand points and discrete inventory levels is prohibitive for them to be pre-computed. However, under Assumption 4.2 of iid error terms from any distribution, the probability information can be partially pre-processed at the same time preserving tractability.

Under Assumption 4.2 of iid error terms, the distribution of \mathcal{E} depends only on the number of error terms summed in expression (C.5), which is bounded by $N = |\mathcal{P}|(|\mathcal{T}| - 1)$ as discussed in Section 4.3.3 of Chapter 4. Pre-processing is performed by choosing a sufficiently large number M and for $m \in \{1, \dots, M\}$ simulating:

$$e_g^m = \sum_{t=1}^g \varepsilon_{it}, \quad \forall g \in \{1, \dots, N\}, \tag{C.6}$$

by drawing ε_{it} from the marginal distribution Φ' for any $i \in \mathcal{P}$ (see Section 4.3.3 in Chapter 4). Using the result of (C.6), we derive an empirical distribution function Φ_g^{emp} of the values $\{e_g^1, \dots, e_g^M\}, \forall g \in \{1, \dots, N\}$. Given Assumption 4.2 and formulation (C.6), $\exists g \in \{1, \dots, N\}$ s.t. $\mathcal{E} \sim \Phi_g^{\text{emp}}$. These empirical distribution functions are then used at runtime to calculate the probabilities in formula (C.4).

Bibliography

- Abdelmaguid, T. F. (2004). *Heuristic approaches for the integrated inventory distribution problem*. PhD thesis, University of Southern California, Los Angeles, CA, USA.
- Abdollahi, M., Arvan, M., Omidvar, A., and Ameri, F. (2014). A simulation optimization approach to apply value at risk analysis on the inventory routing problem with backlogged demand. *International Journal of Industrial Engineering Computations*, 5(4):603–620.
- Abramowitz, M. and Stegun, I. A., editors (1972). *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. New York: Dover Publications.
- Adler, J. D. and Mirchandani, P. B. (2017). The vehicle scheduling problem for fleets with alternative-fuel vehicles. *Transportation Science*, 51(2):441–456.
- Adulyasak, Y., Cordeau, J.-F., and Jans, R. (2015). Benders decomposition for production routing under demand uncertainty. *Operations Research*, 63(4):851–867.
- Aghezzaf, E.-H. (2008). Robust distribution planning for supplier-managed inventory agreements when demand rates and travel times are stationary. *The Journal of the Operational Research Society*, 59(8):1055–1065.
- Andersson, H., Hoff, A., Christiansen, M., Hasle, G., and Løkketangen, A. (2010). Industrial aspects and literature survey: Combined inventory management and routing. *Computers & Operations Research*, 37(9):1515–1536.
- Angelelli, E. and Speranza, M. G. (2002a). The application of a vehicle routing model to a waste-collection problem: Two case studies. *The Journal of the Operational Research Society*, 53(9):944–952.
- Angelelli, E. and Speranza, M. G. (2002b). The periodic vehicle routing problem with intermediate facilities. *European Journal of Operational Research*, 137(2):233–247.
- Anghinolfi, D., Paolucci, M., Robba, M., and Taramasso, A. C. (2013). A dynamic optimization model for solid waste recycling. *Waste Management*, 33(2):287–296.
- Archetti, C., Bertazzi, L., Hertz, A., and Speranza, M. G. (2012). A hybrid heuristic for an inventory routing problem. *INFORMS Journal on Computing*, 24(1):101–116.

Bibliography

- Archetti, C., Bertazzi, L., Laporte, G., and Speranza, M. G. (2007). A branch-and-cut algorithm for a vendor-managed inventory-routing problem. *Transportation Science*, 41(3):382–391.
- Archetti, C., Bertazzi, L., Paletta, G., and Speranza, M. G. (2011). Analysis of the maximum level policy in a production-distribution system. *Computers & Operations Research*, 38(12):1731–1746.
- Arribas, C. A., Blazquez, C. A., and Lamas, A. (2010). Urban solid waste collection system using mathematical modelling and tools of geographic information systems. *Waste Management & Research*, 28(4):355–363.
- Arslan, A. N. and Papageorgiou, D. J. (2015). Bulk ship fleet renewal and deployment under uncertainty: A multi-stage stochastic programming approach. Working Paper, Department of Industrial & Systems Engineering, University of Florida, Gainesville, FL, USA.
- Asimakopoulos, G., Christodoulou, S., Alefragis, P., Gizas, A., and Triantafyllou, V. (2016). Dynacargo: The evaluation results of a dynamic waste collection management system based on real-time and forecasted data. In *Proceedings of the 20th Pan-Hellenic Conference on Informatics*, PCI '16, pages 7:1–7:5, Patras, Greece.
- Baldacci, R. and Mingozzi, A. (2009). A unified exact method for solving different classes of vehicle routing problems. *Mathematical Programming*, 120(2):347–380.
- Bard, J. F., Huang, L., Dror, M., and Jaillet, P. (1998a). A branch and cut algorithm for the VRP with satellite facilities. *IIE Transactions*, 30(9):821–834.
- Bard, J. F., Huang, L., Jaillet, P., and Dror, M. (1998b). A decomposition approach to the inventory routing problem with satellite facilities. *Transportation Science*, 32(2):189–203.
- Beliën, J., Boeck, L. D., and Ackere, J. V. (2014). Municipal solid waste collection and management problems: A literature review. *Transportation Science*, 48(1):78–102.
- Bell, W. J., Dalberto, L. M., Fisher, M. L., Greenfield, A. J., Jaikumar, R., Kedia, P., Mack, R. G., and Prutzman, P. J. (1983). Improving the distribution of industrial gases with an on-line computerized routing and scheduling optimizer. *Interfaces*, 13(6):4–23.
- Beltrami, E. J. and Bodin, L. D. (1974). Networks and vehicle routing for municipal waste collection. *Networks*, 4(1):65–94.
- Benjamin, A. M. (2011). *Metaheuristics for the Waste Collection Vehicle Routing Problem with Time Windows*. PhD thesis, Department of Mathematical Sciences, Brunel University, London, UK.
- Bertazzi, L., Bosco, A., Guerriero, F., and Laganà, D. (2013). A stochastic inventory routing problem with stock-out. *Transportation Research Part C: Emerging Technologies*, 27:89–107.

- Bertazzi, L., Bosco, A., and Laganà, D. (2015). Managing stochastic demand in an inventory routing problem with transportation procurement. *Omega*, 56:112–121.
- Bertazzi, L., Paletta, G., and Speranza, M. G. (2002). Deterministic order-up-to level policies in an inventory routing problem. *Transportation Science*, 36(1):119–132.
- Bertazzi, L. and Speranza, M. (2013). Inventory routing problems with multiple customers. *EURO Journal on Transportation and Logistics*, 2(3):255–275.
- Bertsimas, D. and Sim, M. (2003). Robust discrete optimization and network flows. *Mathematical Programming*, 98(1-3):49–71.
- Bertsimas, D. and Sim, M. (2004). The price of robustness. *Operations Research*, 52(1):35–53.
- Bertsimas, D. J., Jaillet, P., and Odoni, A. R. (1990). A priori optimization. *Operations Research*, 38(6):1019–1033.
- Bettinelli, A., Ceselli, A., and Righini, G. (2011). A branch-and-cut-and-price algorithm for the multi-depot heterogeneous vehicle routing problem with time windows. *Transportation Research Part C: Emerging Technologies*, 19(5):723–740.
- Bitsch, B. (2012). Inventory routing with stochastic demand. Master's thesis, Aarhus School of Business and Social Sciences, Aarhus University, Aarhus, Denmark.
- Buhrkal, K., Larsen, A., and Ropke, S. (2012). The waste collection vehicle routing problem with time windows in a city logistics context. *Procedia-Social and Behavioral Sciences*, 39:241–254.
- Campbell, A. M. and Savelsbergh, M. W. P. (2004). A decomposition approach for the inventory-routing problem. *Transportation Science*, 38(4):488–502.
- Chao, I.-M., Golden, B., and Wasil, E. (1999). A computational study of a new heuristic for the site-dependent vehicle routing problem. *INFOR: Information Systems and Operational Research*, 37(3):319–336.
- Cheng, L. and Duran, M. A. (2004). Logistics for world-wide crude oil transportation using discrete event simulation and optimal control. *Computers & Chemical Engineering*, 28(6–7):897–911.
- Christodoulou, S., Alefragis, P., Gizas, A., Asimakopoulos, G., and Triantafyllou, V. (2016). Dynacargo routing subsystem and its algorithms for efficient urban waste collection. In *11th International Conference on Practice and Theory of Automated Timetabling (PATAT-2016)*, Udine, Italy, August 23-26, 2016.
- Coelho, L. C., Cordeau, J.-F., and Laporte, G. (2012a). Consistency in multi-vehicle inventory-routing. *Transportation Research Part C: Emerging Technologies*, 24:270–287.

Bibliography

- Coelho, L. C., Cordeau, J.-F., and Laporte, G. (2012b). The inventory-routing problem with transshipment. *Computers & Operations Research*, 39(11):2537–2548.
- Coelho, L. C., Cordeau, J.-F., and Laporte, G. (2014a). Heuristics for dynamic and stochastic inventory-routing. *Computers & Operations Research*, 52(A):55–67.
- Coelho, L. C., Cordeau, J.-F., and Laporte, G. (2014b). Thirty years of inventory routing. *Transportation Science*, 48(1):1–19.
- Conrad, R. G. and Figliozzi, M. A. (2011). The recharging vehicle routing problem. In Doolen, T. and Aken, E. V., editors, *Proceedings of the 2011 Industrial Engineering Research Conference*, Reno, NV, USA.
- Cordeau, J.-F., Gendreau, M., and Laporte, G. (1997). A tabu search heuristic for periodic and multi-depot vehicle routing problems. *Networks*, 30(2):105–119.
- Cordeau, J.-F. and Laporte, G. (2001). A tabu search algorithm for the site dependent vehicle routing problem with time windows. *INFOR*, 39(3):292–298.
- Cordeau, J.-F., Laporte, G., and Mercier, A. (2001). A unified tabu search heuristic for vehicle routing problems with time windows. *The Journal of the Operational Research Society*, 52(8):928–936.
- Cordeau, J.-F., Laporte, G., and Mercier, A. (2004). Improved tabu search algorithm for the handling of route duration constraints in vehicle routing problems with time windows. *Journal of the Operational Research Society*, 55:542–546.
- Crevier, B., Cordeau, J.-F., and Laporte, G. (2007). The multi-depot vehicle routing problem with inter-depot routes. *European Journal of Operational Research*, 176(2):756–773.
- Dantzig, G. and Ramser, R. (1959). The truck dispatching problem. *Management Science*, 6:80–91.
- Delage, E. and Iancu, D. A. (2015). Robust multistage decision making. In *The Operations Research Revolution*, INFORMS Tutorials in Operations Research, pages 20–46.
- Dondo, R. and Cerdá, J. (2006). A reactive MILP approach to the multidepot heterogeneous fleet vehicle routing problem with time windows. *International Transactions in Operational Research*, 13(5):441–459.
- Dondo, R. and Cerdá, J. (2007). A cluster-based optimization approach for the multi-depot heterogeneous fleet vehicle routing problem with time windows. *European Journal of Operational Research*, 176(3):1478–1507.
- Dondo, R. G. and Cerdá, J. (2009). A hybrid local improvement algorithm for large-scale multi-depot vehicle routing problems with time windows. *Computers & Chemical Engineering*, 33(2):513–530.

- Dror, M. and Ball, M. (1987). Inventory/routing: Reduction from an annual to a short-period problem. *Naval Research Logistics*, 34(6):891–905.
- Dror, M., Ball, M., and Golden, B. (1985). A computational comparison of algorithms for the inventory routing problem. *Annals of Operations Research*, 4(1):1–23.
- Dror, M. and Levy, L. (1986). A vehicle routing improvement algorithm comparison of a "greedy" and a matching implementation for inventory routing. *Computers & Operations Research*, 13(1):33–45.
- Dror, M. and Trudeau, P. (1986). Stochastic vehicle routing with modified savings algorithm. *European Journal of Operational Research*, 23(2):228–235.
- Dror, M. and Trudeau, P. (1996). Cash flow optimization in delivery scheduling. *European Journal of Operational Research*, 88(3):504–515.
- Erdoğan, S. and Miller-Hooks, E. (2012). A green vehicle routing problem. *Transportation Research Part E: Logistics and Transportation Review*, 48(1):100–114.
- Errorhout, A., Kharraja, S., and Corbier, C. (2016). Two-stage stochastic assignment problem in the home health care. *IFAC-PapersOnLine*, 49(12):1152–1157. 8th IFAC Conference on Manufacturing Modelling, Management and Control MIM, Troyes, France, June 28–30, 2016.
- Errorhout, A., Kharraja, S., and Matta, A. (2014). The uncertainty in the home health care assignment problem. In *Proceedings of the 3rd International Conference on Operations Research and Enterprise Systems*, pages 453–459, Angers, France, March 6–8, 2014.
- European Commission (2016). *Circular Economy Strategy*. <http://ec.europa.eu/environment/circular-economy>, Accessed on April 16, 2016.
- Faccio, M., Persona, A., and Zanin, G. (2011). Waste collection multi objective model with real time traceability data. *Waste Management*, 31(12):2391–2405.
- Garey, M. R. and Johnson, D. S. (1979). *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., New York, NY, USA.
- Gendreau, M., Jabali, O., and Rei, W. (2014). Chapter 8: Stochastic vehicle routing problems. In Toth, P. and Vigo, D., editors, *Vehicle Routing: Problems, Methods, and Applications, Second Edition*. SIAM.
- Gendreau, M., Jabali, O., and Rei, W. (2016). 50th anniversary invited article–Future research directions in stochastic vehicle routing. *Transportation Science*, 50(4):1163–1173.
- Ghose, M., Dikshit, A., and Sharma, S. (2006). A GIS based transportation model for solid waste disposal–A case study on Asansol municipality. *Waste Management*, 26(11):1287–1293.

Bibliography

- Goeke, D. and Schneider, M. (2015). Routing a mixed fleet of electric and conventional vehicles. *European Journal of Operational Research*, 245(1):81–99.
- Golden, B., Assad, A., Levy, L., and Gheysens, F. (1984). The fleet size and mix vehicle routing problem. *Computers & Operations Research*, 11(1):49–66.
- Golden, B. L., Assad, A. A., and Wasil, E. A. (2002). Routing vehicles in the real world: Applications in the solid waste, beverage, food, dairy, and newspaper industries. In Toth, P. and Vigo, D., editors, *The Vehicle Routing Problem*, pages 245–286. SIAM, Philadelphia, PA.
- Gounaris, C. E., Wiesemann, W., and Floudas, C. A. (2013). The robust capacitated vehicle routing problem under demand uncertainty. *Operations Research*, 61(3):677–693.
- Greco, G., Allegrini, M., Lungo, C. D., Savellini, P. G., and Gabellini, L. (2015). Drivers of solid waste collection costs. Empirical evidence from Italy. *Journal of Cleaner Production*, 106:364–371.
- Halvorsen-Weare, E. E., Fagerholt, K., and Rönnqvist, M. (2013). Vessel routing and scheduling under uncertainty in the liquefied natural gas business. *Computers & Industrial Engineering*, 64(1):290–301.
- Hemmelmayr, V., Doerner, K. F., Hartl, R. F., and Rath, S. (2013). A heuristic solution method for node routing based solid waste collection problems. *Journal of Heuristics*, 19(2):129–156.
- Hemmelmayr, V., Doerner, K. F., Hartl, R. F., and Savelsbergh, M. W. (2010). Vendor managed inventory for environments with stochastic product usage. *European Journal of Operational Research*, 202(3):686–695.
- Hemmelmayr, V. C., Doerner, K. F., Hartl, R. F., and Vigo, D. (2014). Models and algorithms for the integrated planning of bin allocation and vehicle routing in solid waste management. *Transportation Science*, 48(1):103–120.
- Hiermann, G., Puchinger, J., Ropke, S., and Hartl, R. F. (2016). The electric fleet size and mix vehicle routing problem with time windows and recharging stations. *European Journal of Operational Research*, 252(3):995–1018.
- Ivarsøy, K. S. and Solhaug, I. E. (2014). Optimization of combined ship routing and inventory management in the salmon farming industry. Master’s thesis, Norwegian Institute of Science and Technology, Trondheim, Norway.
- Jaillet, P., Bard, J. F., Huang, L., and Dror, M. (2002). Delivery cost approximations for inventory routing problems in a rolling horizon framework. *Transportation Science*, 36(3):292–300.
- Johansson, O. M. (2006). The effect of dynamic scheduling and routing in a solid waste management system. *Waste Management*, 26(8):875–885.

- Kim, B. I., Kim, S., and Sahoo, S. (2006). Waste collection vehicle routing problem with time windows. *Computers & Operations Research*, 33(12):3624–3642.
- Koç, Ç. (2016). A unified-adaptive large neighborhood search metaheuristic for periodic location-routing problems. *Transportation Research Part C: Emerging Technologies*, 68:265–284.
- Krikke, H., le Blanc, I., van Krieken, M., and Fleuren, H. (2008). Low-frequency collection of materials disassembled from end-of-life vehicles: On the value of on-line monitoring in optimizing route planning. *International Journal of Production Economics*, 111(2):209–228.
- Kritzinger, S., Tricoire, F., Doerner, K. F., Hartl, R. E., and Stützle, T. (2017). A unified framework for routing problems with a fixed fleet size. *International Journal of Metaheuristics*, Forthcoming.
- Lahyani, R., Khemakhem, M., and Semet, F. (2015). Rich vehicle routing problems: From a taxonomy to a definition. *European Journal of Operational Research*, 241(1):1–14.
- Lanzarone, E. and Matta, A. (2009). Value of perfect information in home care human resource planning with continuity of care. In *Proceedings of the 35th Conference on Operational Research Applied to Health Services (ORAHS 2009)*, Leuven, Belgium, July 13–17, 2009.
- Lanzarone, E. and Matta, A. (2012). A cost assignment policy for home care patients. *Flexible Services and Manufacturing Journal*, 24(4):465–495.
- Lanzarone, E., Matta, A., and Sahin, E. (2012). Operations management applied to home care services: The problem of assigning human resources to patients. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, 42(6):1346–1363.
- Larson, R. C. (1988). Transporting sludge to the 106-mile site: An inventory/routing model for fleet sizing and logistics system design. *Transportation Science*, 22(3):186–198.
- Louveaux, F. (1998). An introduction to stochastic transportation models. In Labbé, M., Laporte, G., Tanczos, K., and Toint, P., editors, *Operations Research and Decision Aid Methodologies in Traffic and Transportation Management*, pages 244–263. Springer Berlin Heidelberg.
- Mancini, S. (2016). A real-life multi depot multi period vehicle routing problem with a heterogeneous fleet: Formulation and adaptive large neighborhood search based matheuristic. *Transportation Research Part C: Emerging Technologies*, 70:100–112.
- Markov, I., de Lapparent, M., Bierlaire, M., and Varone, S. (2015). Modeling a waste disposal process via a discrete mixture of count data models. In *Proceedings of the 15th Swiss Transport Research Conference (STRC)*, April 17–19, 2015, Ascona, Switzerland.

Bibliography

- Markov, I., Varone, S., and Bierlaire, M. (2014). Vehicle routing for a complex waste collection problem. In *Proceedings of the 14th Swiss Transport Research Conference (STRC)*, Ascona, Switzerland.
- McLeod, F., Erdoğan, G., Cherrett, T., Bektaş, T., Davies, N., Speed, C., Dickinson, J., and Norgate, S. (2013). Dynamic collection scheduling using remote asset monitoring. *Transportation Research Record*, 2378:65–72.
- Mendoza, J. E., Castanier, B., Guéret, C., Medaglia, A. L., and Velasco, N. (2011). Constructive heuristics for the multicompartment vehicle routing problem with stochastic demands. *Transportation Science*, 45(3):346–363.
- Mes, M. (2012). Using simulation to assess the opportunities of dynamic waste collection. In Bangsow, S., editor, *Use Cases of Discrete Event Simulation*, pages 277–307. Springer Berlin Heidelberg.
- Mes, M., Schutten, M., and Rivera, A. P. (2014). Inventory routing for dynamic waste collection. *Waste Management*, 34(9):1564–1576.
- Minoux, M. (2009). Robust linear programming with right-hand-side uncertainty, duality and applications. In Floudas, C. A. and Pardalos, P. M., editors, *Encyclopedia of Optimization, 2nd edition*, pages 3317–3327. Springer.
- Moghaddam, N. M. (2015). The partially rechargeable electric vehicle routing problem with time windows and capacitated charging stations. Master's thesis, Clemson University, Clemson, SC, USA.
- Moin, N. H. and Salhi, S. (2007). Inventory routing problems: A logistical overview. *The Journal of the Operational Research Society*, 58(9):1185–1194.
- Moraes, L. A. and Faria, L. F. (2016). A stochastic programming approach to liquefied natural gas planning. *Pesquisa Operacional*, 36:151–165.
- Muter, I., Cordeau, J.-F., and Laporte, G. (2014). A branch-and-price algorithm for the multidepot vehicle routing problem with interdepot routes. *Transportation Science*, 48(3):425–441.
- Nag, B., Golden, B. L., and Assad, A. (1988). Vehicle routing with site dependencies. In Golden, B. L. and Assad, A., editors, *Vehicle Routing: Methods and Studies*, pages 149–159. Elsevier, Amsterdam, The Netherlands.
- Nekooghadirli, N., Tavakkoli-Moghaddam, R., Ghezavati, V., and Javanmard, S. (2014a). Solving a new bi-objective location-routing-inventory problem in a distribution network by meta-heuristics. *Computers & Industrial Engineering*, 76:204–221.
- Nekooghadirli, N., Tavakkoli-Moghaddam, R., and Ghezavati, V. R. (2014b). Efficiency of a multi-objective imperialist competitive algorithm: A biobjective location-routing-inventory problem with probabilistic routes. *Journal of AI and Data Mining*, 2(2):105–112.

- Niakan, F. and Rahimi, M. (2015). A multi-objective healthcare inventory routing problem; a fuzzy possibilistic approach. *Transportation Research Part E: Logistics and Transportation Review*, 80:74–94.
- Nolz, P. C., Absi, N., and Feillet, D. (2011). Optimization of infectious medical waste collection using RFID. In Böse, J. W., Hu, H., Jahn, C., Shi, X., Stahlbock, R., and Voß, S., editors, *Computational Logistics*, volume 6971 of *Lecture Notes in Computer Science*, pages 86–100. Springer Berlin Heidelberg.
- Nolz, P. C., Absi, N., and Feillet, D. (2014a). A bi-objective inventory routing problem for sustainable waste management under uncertainty. *Journal of Multi-Criteria Decision Analysis*, 21(5-6):299–314.
- Nolz, P. C., Absi, N., and Feillet, D. (2014b). A stochastic inventory routing problem for infectious medical waste collection. *Networks*, 63(1):82–95.
- Oliveira Simonetto, E. and Borenstein, D. (2007). A decision support system for the operational planning of solid waste collection. *Waste Management*, 27(10):1286–1297.
- Ombuki-Berman, B. M., Runka, A., and Hanshar, F. T. (2007). Waste collection vehicle routing problem with time windows using multi-objective genetic algorithms. Technical Report CS-07-04, Brock University, St. Catharines, ON, Canada.
- Papageorgiou, D. J., Nemhauser, G. L., Sokol, J., Cheon, M.-S., and Keha, A. B. (2014). MIRPLib—A library of maritime inventory routing problem instances: Survey, core model, and benchmark results. *European Journal of Operational Research*, 235(2):350–366.
- Park, Y.-B., Yoo, J.-S., and Park, H.-S. (2016). A genetic algorithm for the vendor-managed inventory routing problem with lost sales. *Expert Systems with Applications*, 53:149–159.
- Pelletier, S., Jabali, O., and Laporte, G. (2016). 50th anniversary invited article—Goods distribution with electric vehicles: Review and research perspectives. *Transportation Science*, 50(1):3–22.
- Penna, P. H. V., Subramanian, A., and Ochi, L. S. (2013). An iterated local search heuristic for the heterogeneous fleet vehicle routing problem. *Journal of Heuristics*, 19(2):201–232.
- Pillac, V., Gendreau, M., Guéret, C., and Medaglia, A. L. (2013). A review of dynamic vehicle routing problems. *European Journal of Operational Research*, 225(1):1–11.
- Pisinger, D. and Ropke, S. (2007). A general heuristic for vehicle routing problems. *Computers & Operations Research*, 34(8):2403–2435.
- Poggi, M. and Uchoa, E. (2014). Chapter 3: New exact algorithms for the capacitated vehicle routing problem. In Toth, P. and Vigo, D., editors, *Vehicle Routing: Problems, Methods, and Applications, Second Edition*. SIAM.

Bibliography

- Polacek, M., Hartl, R. F., Doerner, K., and Reimann, M. (2004). A variable neighborhood search for the multi depot vehicle routing problem with time windows. *Journal of Heuristics*, 10(6):613–627.
- Powell, W. B. (2011). *Approximate Dynamic Programming: Solving the Curses of Dimensionality, Second Edition*. John Wiley & Sons.
- Rahimi, M., Baboli, A., and Rezik, Y. (2017). Multi-objective inventory routing problem: A stochastic model to consider profit, service level and green criteria. *Transportation Research Part E: Logistics and Transportation Review*, 101:59–83.
- Repoussis, P., Paraskevopoulos, D., Zobelos, G., Tarantilis, C., and Ioannou, G. (2009). A web-based decision support system for waste lube oils collection and recycling. *European Journal of Operational Research*, 195(3):676–700.
- Repoussis, P., Tarantilis, C., Bräysy, O., and Ioannou, G. (2010). A hybrid evolution strategy for the open vehicle routing problem. *Computers & Operations Research*, 37(3):443–455.
- Repoussis, P. P., Tarantilis, C. D., and Ioannou, G. (2007). The open vehicle routing problem with time windows. *Journal of the Operational Research Society*, 58:355–367.
- Ribeiro, R. and Lourenço, H. (2003). Inventory-routing model for a multi-period problem with stochastic and deterministic demand. Technical Report 275, Department of Economics and Business, Universitat Pompeu Fabra, Barcelona, Spain.
- Ritzinger, U., Puchinger, J., and Hartl, R. F. (2016). A survey on dynamic and stochastic vehicle routing problems. *International Journal of Production Research*, 54(1):215–231.
- Rochat, Y. and Taillard, É. D. (1995). Probabilistic diversification and intensification in local search for vehicle routing. *Journal of Heuristics*, 1:147–167.
- Ropke, S. and Pisinger, D. (2006a). An adaptive large neighborhood search heuristic for the pickup and delivery problem with time windows. *Transportation Science*, 40(4):455–472.
- Ropke, S. and Pisinger, D. (2006b). A unified heuristic for a large class of vehicle routing problems with backhauls. *European Journal of Operational Research*, 171(3):750–775.
- Rossi, R., Tomasella, M., Martin-Barragan, B., Embley, T., Walsh, C., and Langston, M. (2017). The dynamic bowser routing problem. Technical report, Business School, University of Edinburgh, UK.
- Rovetta, A., Xiumin, F., Vicentini, F., Minghua, Z., Giusti, A., and Qichang, H. (2009). Early detection and evaluation of waste through sensorized containers for a collection monitoring application. *Waste Management*, 29(12):2939–2949.
- Sahoo, S., Kim, S., Kim, B.-I., Kraas, B., and Popov, A. (2005). Routing optimization for waste management. *Interfaces*, 35(1):24–36.

- Salavati-Khoshghalb, M., Gendreau, M., Jabali, O., and Rei, W. (2017). A rule-based recourse for the vehicle routing problem with stochastic demands. Technical report, CIRRELT-2017-36, Montreal, Canada.
- Sariklis, D. and Powell, S. (2000). A heuristic method for the open vehicle routing problem. *The Journal of the Operational Research Society*, 51(5):564–573.
- Sassi, O., Cherif, W. R., and Oulamara, A. (2014). Vehicle routing problem with mixed fleet of conventional and heterogeneous electric vehicles and time dependent charging costs, Working paper, University of Lorraine, Nancy, France.
- Savelsbergh, M. W. P. (1992). The vehicle routing problem with time windows: Minimizing route duration. *INFORMS Journal on Computing*, 4(2):146–154.
- Schneider, M., Stenger, A., and Goeke, D. (2014). The electric vehicle-routing problem with time windows and recharging stations. *Transportation Science*, 48(4):500–520.
- Schneider, M., Stenger, A., and Hof, J. (2015). An adaptive VNS algorithm for vehicle routing problems with intermediate stops. *OR Spectrum*, 37(2):353–387.
- Semet, F., Toth, P., and Vigo, D. (2014). Chapter 2: Classical exact algorithms for the capacitated vehicle routing problem. In Toth, P. and Vigo, D., editors, *Vehicle Routing: Problems, Methods, and Applications, Second Edition*. SIAM.
- Shaw, P. (1997). A new local search algorithm providing high quality solutions to vehicle routing problems. Technical report, APES Group, Department of Computer Sciences, University of Strathclyde, Glasgow, Scotland.
- Shi, Y., Boudouh, T., and Grunder, O. (2017). A hybrid genetic algorithm for a home health care routing problem with time window and fuzzy demand. *Expert Systems with Applications*, 72:160–176.
- SITG (2017). Points de collecte des déchets. In *Système d'Information du Territoire à Genève*. http://ge.ch/sitg/geodata/SITG/OPENDATA/8438/SHP_GOL_DECHETTERIE.zip, Accessed on February 8, 2017.
- Solomon, M. M. (1987). Algorithms for the vehicle routing and scheduling problems with time window constraints. *Operations Research*, 35(2):254–265.
- Solyali, O., Cordeau, J.-F., and Laporte, G. (2012). Robust inventory routing under demand uncertainty. *Transportation Science*, 46(3):327–340.
- Soysal, M. (2016). Closed-loop inventory routing problem for returnable transport items. *Transportation Research Part D: Transport and Environment*, 48:31–45.
- Soysal, M., Bloemhof-Ruwaard, J. M., Haijema, R., and van der Vorst, J. G. (2015). Modeling an inventory routing problem for perishable products with environmental considerations and demand uncertainty. *International Journal of Production Economics*, 164:118–133.

Bibliography

- Soysal, M., Bloemhof-Ruwaard, J. M., Haijema, R., and van der Vorst, J. G. (2016). Modeling a green inventory routing problem for perishable products with horizontal collaboration. *Computers & Operations Research*. Available online.
- Soyster, A. L. (1973). Technical note—Convex programming with set-inclusive constraints and applications to inexact linear programming. *Operations Research*, 21(5):1154–1157.
- Stewart, W. R. and Golden, B. L. (1983). Stochastic vehicle routing: A comprehensive approach. *European Journal of Operational Research*, 14(4):371–385.
- Subramanian, A., Penna, P. H. V., Uchoa, E., and Ochi, L. S. (2012). A hybrid algorithm for the heterogeneous fleet vehicle routing problem. *European Journal of Operational Research*, 221(2):285–295.
- Sungur, I., Ordóñez, F., and Dessouky, M. (2008). A robust optimization approach for the capacitated vehicle routing problem with demand uncertainty. *IIE Transactions*, 40(5):509–523.
- Taillard, É. D. (1999). A heuristic column generation method for the heterogeneous fleet VRP. *RAIRO - Operations Research*, 33(1):1–14.
- Tarantilis, C. D., Zachariadis, E. E., and Kiranoudis, C. T. (2008). A hybrid guided local search for the vehicle-routing problem with intermediate replenishment facilities. *INFORMS Journal on Computing*, 20(1):154–168.
- Tavares, G., Zsigraiova, Z., Semiao, V., and Carvalho, M. (2009). Optimisation of MSW collection routes for minimum fuel consumption using 3D GIS modelling. *Waste Management*, 29:1176–1185.
- Trudeau, P. and Dror, M. (1992). Stochastic inventory routing: Route design with stockouts and route failures. *Transportation Science*, 26(3):171–184.
- Vidal, T., Crainic, T. G., Gendreau, M., and Prins, C. (2014). A unified solution framework for multi-attribute vehicle routing problems. *European Journal of Operational Research*, 234(3):658–673.
- Yavuz, M. and Çapar, İ. (2017). Alternative-fuel vehicle adoption in service fleets: Impact evaluation through optimization modeling. *Transportation Science*, 51(2):480–493.
- Yousefikhoshbakht, M., Didehvar, F., and Rahmati, F. (2014). Solving the heterogeneous fixed fleet open vehicle routing problem by a combined metaheuristic algorithm. *International Journal of Production Research*, 52(9):2565–2575.
- Yu, Y. (2009). *Stochastic ship fleet routing with inventory Limits*. PhD thesis, University of Edinburgh, UK.

- Yu, Y., Chu, C., Chen, H., and Chu, F. (2012). Large scale stochastic inventory routing problems with split delivery and service level constraints. *Annals of Operations Research*, 197(1):135–158.
- Zamorano, M., Molero, E., Grindlay, A., Rodríguez, M., Hurtado, A., and Calvo, F. (2009). A planning scenario for the application of geographical information systems in municipal waste collection: A case of Churriana de la Vega (Granada, Spain). *Resources, Conservation and Recycling*, 54:123–133.
- Zhalechian, M., Tavakkoli-Moghaddam, R., Zahiri, B., and Mohammadi, M. (2016). Sustainable design of a closed-loop location-routing-inventory supply chain network under mixed uncertainty. *Transportation Research Part E: Logistics and Transportation Review*, 89:182–214.
- Zheng, S. and Chen, S. (2016). Fleet replacement decisions under demand and fuel price uncertainties. *Transportation Research Part D: Transport and Environment*. Available online.

curriculum vitae for

ILIYA MARKOV

Chemin de Chandieu 22, 1006 Lausanne, Switzerland

+41 79 886 52 53 – iliya.markov@epfl.ch

ABOUT ME

I completed my PhD at the Transport and Mobility Laboratory (Transp-OR) of École Polytechnique Fédérale de Lausanne (EPFL) under the supervision of Prof. Michel Bierlaire. In my research, I develop models, and design and implement optimization algorithms for complex transportation systems, in particular vehicle and inventory routing problems.

EDUCATION

École Polytechnique Fédérale de Lausanne (EPFL), Switzerland	<i>Sep, 2013</i>
Doctoral Program in Civil and Environmental Engineering	<i>– Nov, 2017</i>
University of Edinburgh, Scotland	<i>Sep, 2009</i>
Operational Research with Finance, MSc with Distinction	<i>– Nov, 2010</i>
American University in Bulgaria	<i>Sep, 2005</i>
Mathematics and Economics, BA Magna cum Laude	<i>– May, 2009</i>

WORK EXPERIENCE

École Polytechnique Fédérale de Lausanne (EPFL), Switzerland	
· <i>Research and teaching assistant:</i> My research is focused on modeling and solving real-world vehicle and inventory routing problems, integrating stochastic demand and forecasting techniques. I have published in international journals and presented at numerous conferences in Switzerland and abroad. I am also involved in the OR-related courses given by the lab.	<i>May, 2016</i> <i>– present</i>
Haute Ecole de Gestion de Genève (HEG), Switzerland	
· <i>Research assistant:</i> I was the main scientific developer and programmer of a decision support tool for solving vehicle and inventory routing problems faced by several recyclable waste collectors.	<i>Oct, 2013</i> <i>– Apr, 2016</i>
· <i>Teaching assistant:</i> I collaborated internationally in publishing two articles in the domain of finance with an application of OR techniques. I also developed a decision support tool for planning the problem of cleaning and moving containers faced by several waste collectors. I was additionally involved in finance and OR-related courses.	<i>May 2011</i> <i>– Sep, 2013</i>

PROJECTS

Ecological Waste Management	
· This CTI project was the focus of my PhD research. I solved numerous scientific challenges related to the modeling and solving of vehicle and inventory routing problems with numerous practice-driven constraints, and the integration of demand uncertainty and forecasting. The project involved close collaboration with our industrial partner EcoWaste SA.	<i>Oct, 2013</i> <i>– Mar, 2016</i>

Optimization of Waste Container Displacements and Cleaning

- I adapted a dial-a-ride genetic algorithm for the container displacements and cleaning problem faced by EcoWaste SA, which I integrated in an interactive graphical user interface for defining and solving the problem, plotting the results, and exporting them in a GPS readable format. Feb, 2013
– Aug, 2013

RESEARCH INTERNSHIPS

CIRRELT, Montreal, Canada

- My research stay, supervised by Prof. Jean-François Cordeau, resulted in the formulation of a real-world stochastic inventory routing problem and the development of a solution methodology. I presented this work at several scientific conferences in 2016, and we co-authored an article. Oct, 2015
– Dec, 2015

SELECTED PUBLICATIONS

Papers in International Journals

- Scarinci, R., Markov, I., and Bierlaire, M. (2017). Network design of a transport system based on accelerating moving walkways, *Transportation Research Part C: Emerging Technologies* 80: 310–328.
- Markov, I., Varone, S., and Bierlaire, M. (2016). Integrating a heterogeneous fixed fleet and a flexible assignment of destination depots in the waste collection VRP with intermediate facilities. *Transportation Research Part B: Methodological*, 84: 256–273.
- Dewaele, B., Markov, I., Pirotte, H., and Tuchschnid, N.S. (2013). Does manager offshore experience count in the alternative UCITS universe? *Journal of Alternative Investments*, 16(1): 72–85.
- Markov, I., Ouevray, R., and Tuchschnid, N.S. (2013). Non-fully invested derivative-free bond index replication. *Financial Markets and Portfolio Management*, 27(1): 101–124.

Papers in Conference Proceedings

- Markov, I., Bierlaire, M., Cordeau, J.-F., Maknoon, Y., and Varone, S. (2017). A general framework for routing problems with stochastic demands. *Proceedings of the 17th Swiss Transport Research Conference (STRC)*, May, 17-19, 2017.
- Markov, I., Lapparent, M. (de), Bierlaire, M., and Varone, S. (2015). Modeling a waste disposal process via a discrete mixture of count data models. *Proceedings of the 15th Swiss Transport Research Conference (STRC)*, April 15-17, 2015, Ascona, Switzerland.
- Markov, I., Varone, S., and Bierlaire, M. (2014). Vehicle routing for a complex waste collection problem. *Proceedings of the 14th Swiss Transport Research Conference (STRC)*, May 14-16, 2014, Ascona, Switzerland.

Technical Reports

- Scarinci, R., Markov, I., and Bierlaire, M. (2017). Network design of a transport system based on accelerating moving walkways. Technical report TRANSP-OR 170123, Transport and Mobility Laboratory, EPFL, Lausanne, Switzerland.
- Markov, I., Bierlaire, M., Cordeau, J.-F., Maknoon, Y., and Varone, S. (2016). Inventory routing with non-stationary stochastic demands. Technical report TRANSP-OR 160825, Transport and Mobility Laboratory, EPFL, Lausanne, Switzerland.

- Markov, I., Varone, S., and Bierlaire, M. (2015). The waste collection VRP with intermediate facilities, a heterogeneous fixed fleet and a flexible assignment of origin and destination depot. Technical report TRANSP-OR 150212, Transport and Mobility Laboratory, EPFL, Lausanne, Switzerland.

TEACHING EXPERIENCE

Teaching Assistantship

Optimization and Simulation Spring 2016
EPFL, Doctoral level, Exercises and labs

Decision-aid Methodologies in Transportation Spring 2014, 2015, 2017
EPFL, Master level, Exercises and labs

Introduction to Optimization Fall 2014, 2015, 2016, 2017
EPFL, Bachelor level, Exercises and labs

Probability and Statistics Fall 2013
EPFL, Bachelor level, Exercises and labs

Supervision of Master's Theses

- Raphaël Lüthi, *Analysis, Tuning, Hybridization and Design of Experiments for an Adaptive Large Neighborhood Search (ALNS) Heuristic for a Complex Inventory Routing Problem*, June 24, 2016.
- Bastien Rojanawisut, *Optimisation of the Network Design of a Futuristic Transport System Based on Accelerated Moving Walkways*, January 16, 2015.

Supervision of Semester Projects

- Romain Olivier Bondo Konde, *Models for Pedestrian Movements Based on Integrated and Sequential Clustering*, June 02, 2017.
- Prisca Aeby, *Generalization and Policy Analysis for a Rich Inventory Routing Problem*, December 23, 2016.
- Raphaël Lüthi, *Implementation of a Futuristic Transport System Based on Accelerated Moving Walkways: Optimization on a Real Case Study*, May 29, 2015.
- Thomas Cibils, *Routing of a Mixed Fleet of Electric and Internal Combustion Trucks*, December 19, 2014.

REVIEWING

- EURO Journal on Transportation and Logistics
- Journal of Heuristics
- Transportation Research Part B: Methodological

ACADEMIC AWARDS

- Green Transportation Award at the Triennial Symposium on Transportation Analysis (TRI-STAN IX), June 16, 2016, Oranjestad, Aruba

- Full Postgraduate Scholarship for Operational Research at the University of Edinburgh awarded by the Student Awards Agency for Scotland (SAAS)
- Certificate for Outstanding Achievement in Mathematics at the American University in Bulgaria
- American Foundation for Bulgaria Distinguished Scholarship at the American University in Bulgaria

COMPUTER SKILLS

General-purpose languages	Java, C#, C++, VBA
Scientific computing	R, Matlab, Mathematica
Optimization	Gurobi, CPLEX, GLPK, AMPL
Databases	MySQL, PostgreSQL
GIS and related	PostGIS, OpenLayers, OsmSharp
Web development	PHP, HTML, CSS, JavaScript, jQuery, Bootstrap
General	Windows, Linux, L ^A T _E X, Office

LANGUAGES

Bulgarian	native language	
English	written: proficient	spoken: proficient
French	written: intermediate	spoken: intermediate

CONFERENCES AND SEMINARS

- Markov, I., Bierlaire, M., Cordeau, J.-F. and Varone, S. *A unified modeling and solution framework for stochastic routing problems*. hEART 2017 – 6th Symposium of the European Association for Research in Transportation, September 14, 2017, Technion – Israel Institute of Technology, Haifa, Israel.
- Markov, I., Bierlaire, M., Cordeau, J.-F, Maknoon, Y., and Varone, S. *A unified framework for rich routing problems with stochastic demands*. VeRoLog 2017 – Sixth meeting of the EURO Working Group on Vehicle Routing and Logistics Optimization, July 11, 2017, Vrije Universiteit Amsterdam, The Netherlands.
- Markov, I., Bierlaire, M., Cordeau, J.-F, Maknoon, Y., and Varone, S. *A general framework for routing problems with stochastic demands*. 17th Swiss Transport Research Conference (STRC), May 17, 2017, Ascona, Switzerland.
- Markov, I., Bierlaire, M., Cordeau, J.-F, Maknoon, Y., and Varone, S. *Waste collection inventory routing with non-stationary stochastic demands*. hEART 2016 – 5th Symposium of the European Association for Research in Transportation, September 16, 2016, TU Delft, The Netherlands.
- Markov, I., Bierlaire, M., Cordeau, J.-F, Maknoon, Y., and Varone, S. *Integrating demand uncertainty in inventory routing for recyclable waste collection*. 14th Joint Operations Research Days, September 09, 2016, Università della Svizzera italiana, Lugano, Switzerland.
- Markov, I., Maknoon, Y., Cordeau J.-F, Varone, S., and Bierlaire M. (2016). *Waste collection inventory routing with non-stationary stochastic demands*. TRISTAN IX – The Triennial Symposium on Transportation Analysis, June 17, 2016, Oranjestad, Aruba.

- Markov, I., Maknoon, Y., Cordeau J.-F., Varone, S., and Bierlaire M. (2016). *Waste collection inventory routing with non-stationary stochastic demands*. VeRoLog 2016 – Fifth meeting of the EURO Working Group on Vehicle Routing and Logistics Optimization, June 6, 2016, École des Mines de Nantes, France.
- Markov, I., Maknoon, Y., Cordeau J.-F., Varone, S., and Bierlaire M. (2016). *Inventory routing with non-stationary stochastic demands*. 16th Swiss Transport Research Conference (STRC), May 19, 2016, Ascona, Switzerland.
- Markov, I., Maknoon, Y., Lapparent, M. (de), Cordeau J.-F., Varone, S., and Bierlaire M. (2015). *Vehicle routing and demand forecasting in recyclable waste collection*. CIRRELT Seminar, November 26, 2015, Centre interuniversitaire de recherche sur les réseaux d'entreprise, la logistique et le transport, Montreal, Canada.
- Markov, I., Lapparent, M. (de), Varone, S., and Bierlaire, M. (2015). *Towards an integrated approach for demand forecasting and vehicle routing in recyclable waste collection*. hEART 2015 – 4th Symposium of the European Association for Research in Transportation, September 11, 2015, DTU, Copenhagen, Denmark.
- Markov, I., Lapparent, M. (de), Varone, S., and Bierlaire, M. (2015). *Vehicle routing and demand forecasting in a generalized waste collection problem*. VeRoLog 2015 – Fourth meeting of the EURO Working Group on Vehicle Routing and Logistics Optimization, June 9, 2015, University of Vienna, Austria.
- Markov, I., Lapparent, M. (de), Bierlaire, M., and Varone, S. (2015). *Modeling a waste disposal process via a discrete mixture of count data models*. 15th Swiss Transport Research Conference (STRC), April 16, 2015, Ascona, Switzerland.
- Markov, I., Varone, S., and Bierlaire, M. (2014). *Solving a complex waste collection routing problem with intermediate disposals*. hEART 2014 – 3rd Symposium of the European Association for Research in Transportation, September 12, 2014, University of Leeds, UK.
- Markov, I., Varone, S., and Bierlaire, M. (2014). *Vehicle routing for a complex waste collection problem*. VeRoLog 2014 – Third meeting of the EURO Working Group on Vehicle Routing and Logistics Optimization, June 23, 2014, University of Oslo, Norway.
- Markov, I., Varone, S., and Bierlaire, M. *Vehicle routing for a complex waste collection problem*. 14th Swiss Transport Research Conference (STRC), May 14, 2014, Ascona, Switzerland.
- Markov, I., Varone, S., and Bierlaire, M. (2014). *Waste collection routing with time windows and intermediate disposal trips*. 12th Joint Operations Research Days, May 9, 2014, EPFL, Lausanne, Switzerland.

