The Impact of Charging Electric Vehicles Between Routing Instances

The Modeling of a Combined Electric Vehicle Routing and Charging Model

Y.M. den Bakker

Photo by Sylvain Glim on Unsplash

TU Delft
ORTEC

The Impact of Charging Electric Vehicles Between Routing Instances

The Modeling of a Combined Electric Vehicle Routing and Charging Model

Thesis report

by

Y.M. den Bakker

to obtain the degree of Master of Science
at the Delft University of Technology
to be defended publicly on April 10, 2024 at 09:30

Thesis committee:	
Supervisors:	Dr. K. Aardal
	Dr. A. Bishnoi
External examiner:	Dr. G.F. Nane
Place:	Faculty of Electrical Engineering, Mathematics and Computer Science, Delft
Project Duration:	September, 2023 - April, 2024
Student number:	4725638

An electronic version of this thesis is available at `http://repository.tudelft.nl/`.

Faculty of Electrical Engineering, Mathematics and Computer Science · Delft University of Technology

Acknowledgements

This endeavor would not have been possible without the guidance and support from ORTEC and TU Delft. I am indebted to ORTEC for giving me the opportunity to write my thesis in such a knowledgeable and hospitable environment, and I am grateful for each of my supervisors and colleagues that helped along the way, in particular Lorenzo Simons and Wilco Draijer. Specifically, I want to thank Lorenzo for helping me shape the project to what it has become, offering me guidance on technical and practical matters, and for giving valuable advice. The guidance and aid that Wilco has given me during my time at ORTEC was also incredibly valuable to me, as was his detailed feedback, directly improving the quality of my work.

I would like to extend my sincere thanks to the members of my thesis committee at the university for joining me on this journey. In particular to my supervisor Anurag Bishnoi, whose guidance and trust allowed me to confidently complete my project and reaching its full potential.

I am also grateful to my partner Jesse, who was always there to offer me his opinions and advice, supporting me both emotionally and in the completion of this work. Without him, this project would not have been the same. Lastly, I'd like to acknowledge my family and friends for their support, encouragement and interest in my work, helping me stay motivated and energized.

Yvonne Margarita den Bakker,
The Hague, April 2024

Abstract

E-mobility, in particular electric vehicles (EVs), play a crucial role in the energy transition. While businesses are increasingly adopting EVs, there is still a lot of opportunity to grow. One aspect of this growth is the way these vehicles are used by companies, especially when it comes to the logistics of EV charging. To encourage companies to further reduce their carbon footprint by more efficiently utilizing EVs, this project proposes a combined vehicle routing and charging model. These models can be used separately, as well as together, allowing the charging model to be combined with any pre-existing routing engine. The goal of this project is to show the benefits of allowing EVs to be charged between shifts during the day, rather than exclusively overnight, as well as to show how such schedules can be made. Our results show that giving vehicles the opportunity to charge between shifts can significantly reduce the costs associated with fleet operations. If the fleet contains non-electric vehicles as well as electric ones, we also see a significant reduction in the number of kilometers driven using fossil fuels. When sufficient chargers were available, even when the vehicles had little time to charge, a feasible schedule could always be found. Moreover, when more realistic charging intervals were used, most vehicles were even able to fully recharge before the start of their next shift. Finally, we concluded that the set of chargers needed to find such a feasible schedule can be relatively small, meaning that even without extensive additions to the charging infrastructure, companies can still benefit from this policy change.

Contents

List of Figures ix

List of Tables x

1 Introduction 1

I Literature Review 3

2 Introduction 4

3 Modeling Techniques 5
3.1 Mathematical Modeling Method . . . 5
3.2 Mathematical Problems . . . 10
3.3 Solution Methods . . . 17

4 Battery Modeling 19
4.1 Energy Consumption . . . 19
4.2 Charging . . . 26
4.3 General Model . . . 28

5 Contribution to the Literature 34
5.1 Conclusion . . . 34
5.2 Contributions of this Research . . . 35

II Model Design 36

6 Preliminary Work 37
6.1 Starting Model . . . 37
6.2 Client Data and Interests . . . 38
6.3 Model Objective Difficulties . . . 38
6.4 Our Model Objective . . . 38

7 Modeling Details 40
7.1 Charging Model . . . 40
7.2 Routing Model . . . 45
7.3 Model Variations . . . 54

III Results 59

8 Tests 60
8.1 Baseline Assumptions . . . 60
8.2 Testing Methodology . . . 61
8.3 Experimental Cases . . . 63

9 Results 66
9.1 Case I: VRP . . . 66
9.2 Case II: Charging Schedule + External VRP . . . 68
9.3 Case III: VRP + Charging Schedule . . . 75
9.4 Case IV: Vehicle and Charger Investments . . . 80

IV Closure 83

10 Conclusion 84
10.1 Case I . . . 84
10.2 Case II . . . 84
10.3 Case III . . . 85
10.4 Case IV . . . 85
10.5 Overview . . . 86

11 Discussion 87
11.1 Summary . . . 87
11.2 Interpretations and Implications . . . 87
11.3 Limitations and Considerations . . . 87
11.4 Recommendations . . . 89

References 100

A Metaheuristic Solution Methods for MILPs 101
A.1 Hill Climbing . . . 101
A.2 Variable Neighborhood Search (VNS). . . . 101
A.3 (Adaptive) Large Neighborhood Search ((A)LNS) . . . 101
A.4 Tabu Search . . . 102
A.5 Simulated Annealing . . . 102
A.6 Genetic Algorithms . . . 102
A.7 Ant Colony Optimization . . . 102

B VRP Variations 103

C Heuristic Solution Methods for the EVRP 105
C.1 Common Metaheuristic Approaches. . . . 105
C.2 Hybrid Metaheuristic Approaches . . . 105
C.3 Alternative Approaches . . . 105

Nomenclature

List of Abbreviations

AC	Alternative Current
ALNS	Adaptive Large Neighborhood Search
CC	Constant Current
CEVRP	Capacitated Electric Vehicle Routing Problem
CV	Constant Voltage
DC	Direct Current
DoD	Depth of Discharge
EV	Electric Vehicle
EVRP	Electric Vehicle Routing Problem
EVRPTW	Electric Vehicle Routing Problem with Time Windows
FRVCP	Fixed Route Vehicle Charging Problem
ICE	Internal Combustion Engine
ICEV	Internal Combustion Engine Vehicle
ILP	Integer Linear Programming
IP	Integer Programming
LNS	Large Neighborhood Search
LP	Linear Programming
MILP	Mixed-Integer Linear Programming
MIP	Mixed-Integer Programming
SoC	State of Charge
SOH	State of Health
SWMCP	Self-Weighted Mean Charging Power
TSP	Travelling Salesman Problem
VNS	Variable Neighborhood Search
VRP	Vehicle Routing Problem

Constants

ρ_a	Air density
c_d	Aerodynamic drag coefficient
c_r	Rolling friction coefficient
g	Gravitational constant

List of Symbols

N	Non-Negative integers
Q	Rational numbers
Z	Integers
Δt	Timestep length
$\mathcal{A}'$	Complete set of arcs
$\mathcal{A}$	Set of arcs between each of the customers
$\mathcal{C}'$	Set of customers and (dummy) charging stations
$\mathcal{C}$	Set of customers
$\mathcal{C}_j$	Set of customers and additional nodes j
$\mathcal{H}$	Collection of circuits
$\mathcal{N}$	Set of nodes
ν	Velocity
ν_{min}	Minimal speed for recuperating energy
θ	Gradient angle
ε_G	Energy efficiency of transmission, generator and in-vehicle charger
ε_M	Energy efficiency of transmission, motor and power conversion
a	Acceleration
A_f	Frontal area of the vehicle
$AST_{s,t}$	Binary parameter that determines if shift s is active at timestep t
AVC_j^k	Binary parameter that determines if vehicle k may visit customer j

B Set of battery level intervals

CAP Maximal capacity of the vehicles

$CCI^b_{NL,c}$ Charging speed for interval b for charger c

CCT_c Charger capacity per timestep

CCT_{grid} Total available power grid capacity on each timestep

Ch Set of chargers

COC_j Cost of circuit H_j

CST_{ij} Cost associated with traveling on arc (i, j)

D Set of depots

DEM_j Demand of customer j

E Total energy demand

$e^{\downarrow}$ Braking phase

$e^{\rightarrow}$ Steady speed phase

$e^{\uparrow}$ Acceleration phase

$ECP_{c,t}$ Price per kWh for charger c at timestep t

ECR Energy Consumption Rate of the vehicles per unit of distance

EDV^k Ending depot of vehicle k

EST_i Earliest time to start service at node i

F Set of charging stations

F' Set of dummy charging station nodes

F_a Aerodynamic drag

F_g Gravitational force

F_r Rolling resistance

F_T Traction force

FTA_s Timestep t that shift s starts

H_j Circuit j

i_k Current at timestep k

K Set of vehicles

L Length of a trip

LEN_{ij} Length, or traveling distance, from node i to node j

LST_i Latest time to start service at node i

m Total vehicle mass

MSD^k Maximal duration of shifts driven by vehicle k

MXT Maximal number of trips vehicles may take

P Total power

P_i Total power demand at timestep i

P_M Power demand

P_{acc} Power demand of vehicle accessories

P_{in} Regenerated energy after breaking

P_{out} Total power demand

PER^k Penalty charged when the energy used by partially charged vehicle k exceeds its starting SoC

PMT_i Penalty charged when customer i is not visited by any vehicles

PMW_i Penalty charged when customer i is visited outside of its time window

PST^k_t Binary parameter that determines if vehicle k is driving a previous shift at timestep t

R Internal battery resistance (in chapter 4)

R Set of trips (in chapter 7)

RCK^k Routing cost per kilometer driven by vehicle k

RCR Recharging rate of the batteries

RCT^k Routing cost per task for vehicle k

RCV^k Routing cost of vehicle k

S Set of shifts (in chapter 7)

S Subset of the set of customers $\mathcal{C}$ (in chapter 3)

SDV^k Starting depot of vehicle k

SET_i Service duration at node i

SOC_k Battery level at timestep k

T Duration of a trip (in chapter 3)

T Set of timesteps (in chapter 7)

T' Set of timesteps excluding the final timestep

TDM Total demand

TRT_{ij} Travel time from node i to node j

V_{CV} Terminal voltage upon start of the CV phase

V_{OC} Open voltage of the battery

VBC Vehicle Battery Capacity

VBL^k Battery level of vehicle k before starting its shift

VOR_{ij} Binary parameter determining if vertex i lies on route H_j

YMI_b^k Maximal charge of the battery per section b for vehicle k

YMN_s^k Minimal battery level needed for vehicle k to drive shift s

YMX^k Maximal battery capacity of vehicle k

YRE^k Battery level of vehicle k after returning from the previous shift

List of Figures

3.1 Solution space of problem (3.2). 7
3.2 Solution space of problem (3.2) with added cutting plane. 8
3.3 An example of a branch-and-bound tree. 9

4.1 Speed curve from standing still to a complete stop with steady speed 50 km/h. 21
4.2 The ratio between the average energy use and the energy use over the entire cycle for different vehicle speeds and trip lengths. 25
4.3 Charging curve for AC and DC charging. 28
4.4 Depiction of the CC-CV charging scheme. 29
4.5 Linear approximations in the literature compared to real data. 30
4.6 Real data versus piecewise linear approximation of a $22kW$ charger charging a battery of $16kWh$. 30
4.7 Self-weighted mean charging power for $3.6kW$ charging at different ambient temperatures with and without Battery Thermal Management. 33

7.1 Flowchart of the first model variation. 54
7.2 Flowchart of the second model variation. 55
7.3 Flowchart of the third model variation. 56
7.4 Flowchart of the fourth model variation. 57

9.1 Charging Schedule of time between the second and the first set of routes. 82
9.2 Charging Schedule of time between the third and the second set of routes. 82

11.1 Flowchart of a possible model expansion. 90

List of Tables

4.1 Typically elementary model parameters. 23

9.1 Routing outcomes for the mixed fleet with 50 customers, generated by OHD. 66
9.2 Routing outcomes for the mixed fleet with 50 customers, generated by our routing model. . 67
9.3 Comparison of the outcomes of our routing model and OHD. 67
9.4 Routing outcomes for the fully electric fleet with 50 customers for policy (3) generated by OHD. 67
9.5 Routing outcomes for the fully electric fleet with 50 customers for policy (3) generated by our routing model. 67
9.6 Comparison of the outcomes of our routing model and OHD. 68
9.7 Routing outcomes for the full-sized dataset for policy (1) generated by OHD. 69
9.8 Routing outcomes for the full-sized dataset for policy (3) generated by OHD. 70
9.9 Comparison of the outcomes for policy (1) and policy (3) for the full-sized dataset. 70
9.10 Charging outcomes for the routing pairs of the full-sized dataset under assumption 1. . . . 71
9.11 Charging outcomes for the routing pairs of the full-sized dataset under assumption 2. . . . 71
9.12 Routing outcomes for the mixed fleet with 100 customers for policy (1) generated by OHD. . 72
9.13 Routing outcomes for the mixed fleet with 100 customers for policy (3) generated by OHD. . 72
9.14 Comparison of the outcomes for policy (1) and policy (3) for the mixed fleet with 100 customers. 72
9.15 Comparison of the outcomes for policy (2) and policy (3) for the mixed fleet with 100 customers. 73
9.16 Charging outcomes for the routing pairs of the the mixed fleet with 100 customers under assumption 1. 73
9.17 Charging outcomes for the routing pairs of the mixed fleet with 100 customers under assumption 2. 73
9.18 Routing outcomes for the fully electric fleet with 100 customers for policy (3) generated by OHD. 74
9.19 Comparison of the outcomes for policy (2) and policy (3) for the fully electric fleet with 100 customers. 74
9.20 Charging outcomes for the routing pairs of the fully electric fleet with 100 customers under assumption 1. 75
9.21 Charging outcomes for the routing pairs of the fully electric fleet with 100 customers under assumption 2. 75
9.22 Routing outcomes for the mixed fleet with 50 customers for policy (1) generated by our routing model. 76
9.23 Routing outcomes for the mixed fleet with 50 customers for policy (3) generated by our routing model. 76
9.24 Comparison of the outcomes for policy (1) and policy (3) for the mixed fleet with 50 customers. 76
9.25 Charging outcomes for the routing pairs of the mixed fleet with 50 customers, without post-processing. 77
9.26 Charging outcomes for the routing pairs of the mixed fleet with 50 customers, with post-processing. 77
9.27 Number of vehicles scheduled per charging instance for different modeling scenarios . . . 78
9.28 Comparison of the outcomes for policy (2) and policy (3) for the fully electric fleet with 50 customers. 79
9.29 Charging outcomes for the routing pairs of the fully electric fleet with 50 customers, without post-processing. 79
9.30 Charging outcomes for the routing pairs of the fully electric fleet with 50 customers, with post-processing. 80
9.31 Routing outcomes for the upliftable with 50 customers, generated by OHD. 81
9.32 Routing outcomes for the upliftable with 50 customers, generated by our routing model. . . 81

B.1 Overview of the different VRP variations 104

1 Introduction

In order to reach the goal of limiting global warming to $1.5°C$ agreed upon in the 2015 Paris Agreement [1], avoiding the devastating consequences of more severe climate change impact, it is vital to accelerate the energy transition. Not reaching this goal means that our climate becomes more extreme, risking an increase of severe droughts, heatwaves and rainfall [2]. To make sure that this doesn't happen, we need to significantly reduce the global greenhouse emissions. According to Our World in Data [3], the sector responsible for the second largest amount of greenhouse emissions in 2020 was transport. E-mobility is therefore a crucial part of the energy transition. Electric vehicles (EVs) are a large part of e-mobility. This thesis will focus on the use of EVs. While more and more companies are moving towards carbon-neutral deliveries by means of EVs [4], there is still progress to be made. Not only in the range of the vehicles and the accessible charging infrastructure, two of the largest disadvantages of electric driving, but also in the way these vehicles are used. When more than a few EVs are in use, it becomes harder to efficiently route and charge these vehicles.

A lot of research has already been done about these problems. A literature review about the Electric Vehicle Routing Problem from 2021 [5] identified 136 papers studying this problem and its variations. The problem of charging the vehicles at their depot is often overlooked however; the main focus in these papers lies in charging the EVs on the road. While this problem is most certainly very interesting and widely applicable, not every company wants to take this approach and instead prefers to exclusively charge their vehicles at their depots. When an abundance of chargers is available, the problem of charging all the vehicles overnight is not very challenging. But, when the number of EVs owned is larger than the number of available chargers, this problem becomes more complicated. It might also be desirable to charge vehicles during the day, for example between two shifts. While it might be possible to simply install as many fast-chargers as there are EVs in order to avoid the charging problem, there is a serious price gap between slower AC-charging ports and faster DC-charging ports. Installing fast chargers can easily cost 30 times as much as installing slower chargers, the hardware of such a charger alone can cost as much as a brand-new EV [6]. For that reason, it can really pay off to get creative with limited (fast-)charger availability. Especially when there is only limited time between shifts to charge the vehicles, a simpler and more secure solution is to assume vehicles do not get an opportunity to charge, and only use the range belonging to a single fully charged battery. If it is possible to charge the vehicles during the day, it allows them to drive more kilometers and complete more efficient routes, further increasing the impact caused by replacing conventional vehicles by electric ones.

In this thesis, we design a combined routing and charging model, the goal of which is to create routing and charging schedules allowing the EVs to be used as efficiently as possible. This model consists of a separate routing and charging model, that can function both independently as well as together. The main objective is to show by how much the routes improve if the vehicles driving them are allowed to charge mid-day, instead of only overnight. We perform experiments both with fully electric fleets, and fleets that also contain non-electric vehicles. For the latter, we do not only define success by the number of driven kilometers, but also by the number of kilometers driven by EVs. Another indicator that we use to gauge improvement is the number of EVs that are scheduled. To properly test the capabilities of this model, we set up a few different cases, in which different model configurations and settings are tested. For these tests, we used data from one of ORTEC's clients consisting of a few routing instances. These instances, slightly adapted and/or reduced, are solved by both our routing model, and one of ORTEC's routing models. This allows us to set a baseline on the performance of our model, guaranteeing the quality of our solutions.

This work is divided into four parts. The first part, containing the literature review, is split into four chapters. After a short introduction in chapter 2, we talk about different modeling techniques that are used when solving mathematical optimization problems, in particular electric vehicle routing. The next chapter discusses the way the two different battery processes, charging and discharging, are typically modeled in the context of electric vehicle routing. The last chapter of the literature review summarizes the found results and explains how this work is able to contribute to the literature on this topic. Part II talks about the model design. The first chapter of this part contains the information needed to build the model. The details of this model are then given in the next chapter. Part III contains the part of the project in which the performance of the model is tested an analyzed. We start this section by a chapter explaining the way the tests are set up, and conclude by providing the outcomes of these tests. This leads us to part IV, the closure, in which we draw our conclusions. We end with a discussion evaluating certain aspects of our model and providing some ideas of future research.

Our results show that giving vehicles the opportunity to charge between shifts can significantly reduce the costs associated with the fleet operations. We additionally see a large decrease in the proportion of kilometers that were driven by non-electric vehicles, when those were part of the fleet as well. We also conclude that in order to solve our instances, only a modest set of chargers is needed. This makes it feasible for companies to make this adjustment to their EV planning without the need for extensive and costly investments to their charging infrastructure.

Part I

Literature Review

2
Introduction

In this literature review we discuss the background that will be relevant for this thesis. We start with a chapter on general modeling techniques. This begins with a section on Mixed-Integer Linear Programming (MILP), so that we set a strong basis of this method. Much of the literature on the topic of electric vehicle routing uses MILP formulations to express their models, and we will do the same, so understanding the concept is important. We proceed by describing the different mathematical problems that are used for route planning optimization. The first of these problems is the Vehicle Routing Problem (VRP), which is not yet suitable for Electric Vehicles. This suggests the Electric Vehicle Routing Problem (EVRP) that does take the limited range of EVs into account. There are various modifications of the EVRP that consider additional assumptions or constraints that will also be discussed. The next part is about solution methods. MILP problems can be solved directly using solvers such as CPLEX or Gurobi, but for more complicated problems this might result in very long or even unworkable runtimes. For that reason, we will discuss alternative methods as well.

The next chapter is about modeling the relevant battery-related technologies. When designing a model within this research, it is important to have a clear view of the techniques that are used to accurately model the critical behaviors of EVs: energy consumption and charging. We will describe how the behavior of these processes are modeled, and how routing models incorporate this. The data necessary for an accurate model is also discussed, as well as the impact that variation in different parameters or related external factors have on the performance of the model.

We conclude by providing a summary of these contents and explaining the goal of our research. This is finalized by discussing the academic and practical value of this research and how it contributes to the available literature.

3 Modeling Techniques

This chapter gives a general overview of the ways to model vehicle routing problems. We start off with a section discussing the broad mathematical background. The next two sections discuss the standard vehicle routing problem and its variations, with focus on the variations that apply to EVs. We will also see how these problems generally get solved.

3.1. Mathematical Modeling Method

In this section we will explain what Integer Programming (IP) is, in particular Mixed-Integer Linear Programming, and how it can be used. This includes giving a definition with a small example, mentioning a few common applications, and introducing some of the ways that can be used to solve these types of problems in practice.

3.1.1. Definition

In short, Integer Programming is a branch of mathematical optimization to describe problems of which at least some of the used variables are constrained to take only integer values. If all the constraints of such a problem are linear, we speak of an Integer Linear Program (ILP). To be fully specific, we can define the terms Mixed-Integer Programming (MIP) and Mixed-Integer Linear Programming (MILP) to respectively mean IP and ILP such that at least some of the variables are not restricted to be integer. Many of the problems that we come across in the literature are written as a MILP, or less frequently a MIP. Therefore, it is critical to understand what it means when a problem is of this form.

Any type of IP can be defined by three things: A set of variables and their subsequent domains, a set of constraints, and an objective function. In the case of a MILP, some of these variables must take integer values, while other variables will be allowed to take rational values. Furthermore, the constraints and the objective function will need to be linear. Note that for a MIP this final restriction does not need to hold.

The canonical form of an ILP is the following:

$$\begin{aligned}
&\max_{x} \quad c^T x \\
&\text{s.t.} \\
&Ax \leq b, \\
&\;\; x \geq 0, \\
&\;\; x \in \mathbb{Z}^n
\end{aligned} \tag{3.1}$$

Here, $c^T \in \mathbb{Q}^n$, $b \in \mathbb{Q}^m$ and $A \in \mathbb{Q}^{m \times n}$. This is the most compact way of writing this problem. The i'th row of A, together with the corresponding element b_i signifies a linear constraint. In practice, when presenting a problem of this type, the constraints are written in a more readable manner. A small example of such a MILP is the following:

$$\begin{aligned} \min_{x} \quad & 5x \\ \text{s.t.} \quad & \\ x + 2y &\geq 5, \\ y &\leq 1.4, \\ x, y &\geq 0, \\ x &\in \mathbb{Z} \end{aligned} \tag{3.2}$$

Note that this is a MILP, since we did not restrict the variable y to take integer values. Upon inspection, one could see that the solution to this problem is 15. We only want to minimize the value of x, and since the first constraint tells us that $x + 2y$ is bound from below, we would like to find as large a value of $2y$ as possible. Since the value of y is bound from above, we could set y to be equal to its upper bound 1.4, which leaves the constraint $x \geq 2.2$. This means $x = 3$, since x must be an integer.

Of course, most problems cannot be solved this easily and need much more advanced techniques. Solving problems exactly might not be practically feasible for many practical problems, simply due to the sheer size of the problems: the number of variables or constraints can grow exponentially fast. It is known that Integer Programming is NP-hard [7], which means that while a solution can be checked in polynomial time, it is typically assumed that it cannot be found in polynomial time.

One of the most famous combinatorial problems, that is, problems that aim to find an optimal object from a finite set of objects, is the Traveling Salesman Problem (TSP). It asks for the shortest path to visit a set of cities, where the distance between each city is known, before returning to the origin. This problem can be written as an ILP, and illustrates how these types of problems can grow so quickly. The Dantzig–Fulkerson–Johnson formulation, one of the stronger possible formulations [8], contains so-called subtour elimination constraints [9]. These constraints ensure that the result of the TSP is a single route, and not a collection of smaller routes covering all the cities. In order to do this, every single subtour needs to be eliminated by an individual constraint. This results in a list of constraints that grows exponentially as a function of the number of cities. This problem happens to be polynomially separable [10], which in this case means that we do not need to simultaneously consider this entire set of constraints, and allows us to solve this problem in polynomial time regardless. This is not the case for every problem however, and other methods might be necessary to find solutions.

In the previous example, Integer Programming was used to find a route amongst a set of locations by introducing a decision variable for each possible arc between locations and determining which subset of arcs form a route. An extension of this is the Vehicle Routing Problem, which will be discussed in detail in the next section. In general however, there are many other ways of applying types of Integer Programming. To give a few common examples, decision variables can represent the number of products to be produced under certain conditions, the locations to be partitioned in different territories, or possible time windows for scheduling different time-constrained activities.

3.1.2. Finding Solutions

Generally, MILPs cannot be solved like we did with example problem (3.2). We can however use this problem to illustrate how these problems usually get solved. In figure (3.1) the solution space of this problem is depicted.

The red half-space depicts the first constraint $x + 2y \geq 5$, and the blue half-space depicts the second constraint $y \leq 1.4$. The intersection of the half-spaces that the constraints signify is the region that contains all the feasible points. That means that in this case, any point within the overlap of the blue and red half-spaces such that x is an integer, is a feasible solution. In figure (3.1), all the feasible solutions are depicted by the black dashed lines. We however aren't interested in just any feasible solution, the goal is to find an optimal solution. That brings us to the objective function, in this example min $5x$. The objective function tells us what direction to optimize in. To find this direction, we simply take the gradient of the objective function. Here, we find direction $(5, 0)$, which means we are simply minimizing along the x-axis. This has been depicted in the figure by means of an arrow. The first step to finding an optimal solution is to relax the constraints that fix the decision variable(s) to be integer. The resulting problem is called the Linear Programming (LP) relaxation. Unlike the MILP from before, this problem can often be solved

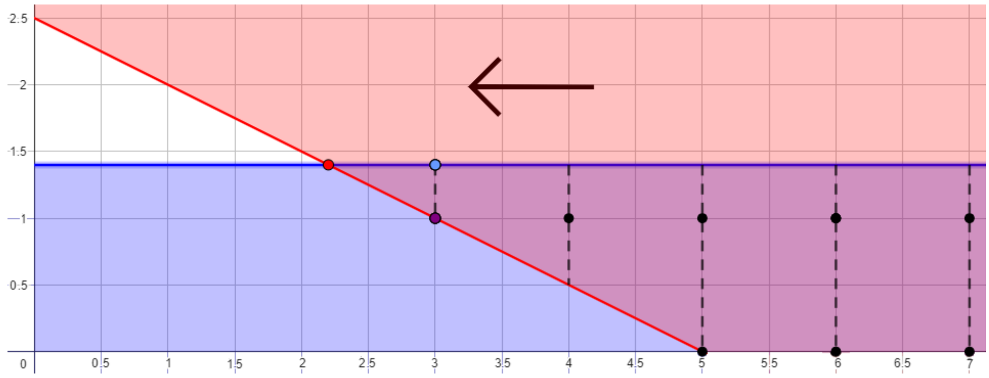

Figure 3.1: Solution space of problem (3.2).

efficiently, using either basic exchange algorithms such as the simplex method, or interior point algorithms, such as the ellipsoid method. For very complicated problems, such as large instances of the TSP, the LP relaxation will still be challenging to solve. There are ways of remedying this, which will be discussed later in this section.

Solving the LP relaxation of a problem often gives infeasible solutions due to relaxing the integer-restriction. They can however already give information on the original problem, since the objective value of the LP relaxation acts as a lower bound for the objective value of the original problem (if it is a minimization problem, otherwise it is an upper bound). In the example we will see that the LP relaxation has an objective value of 12, while the objective value of the MILP is 15. The difference between these two values is called the integrality gap, and represents the quality of the LP relaxation. If this gap is small, we call the LP relaxation strong, as it closely represents the original problem.

In this example, it can be seen immediately that the red point $(2.2, 1.4)$ in figure (3.1) is the optimal value of the LP relaxation, as it is the left-most feasible solution. However, the x-value of this point is not an integer, and therefore not feasible for our MILP. It is not always immediately obvious how to instead find a solution for which x is integer. In this example it is easy to see that the blue point $(3, 1.4)$ is the optimal solution, since it is the nearest feasible solution, but without the picture one might have first tried to round x down to 2, which would have given an infeasible solution. Were y also restricted to be integer, this would have further complicated things, since that point would not have been feasible either. Instead, we find the purple point $(3, 1)$, that is even further away from the solution found by the LP relaxation. There are roughly two categories of methods to algorithmically find the solution to the original MILP: exact methods and heuristics. Both will be discussed below.

It should be noted that there is not always a unique solution. If instead of minimizing a function of x, we had minimized y, we would have seen that there would be many solutions: each integer value above 4 on the x-axis would produce the same objective value of 0. If instead of minimizing a function of x we had maximized this same function, we would have gotten an unbounded problem: the objective value diverges to infinity, since there is no constraint that bounds x from above. Finally, it is also possible that there is no solution at all: this happens when the half-spaces of the constraints do not overlap, or when they do, the overlap does not contain any points satisfying the integer-restrictions. These possibilities further complicate the solution of these kinds of problems.

3.1.3. Exact Methods

The methods of finding exact solutions to (M)ILPs can be divided into two categories: cutting plane methods and variants of the branch-and-bound method.

Cutting Plane Methods

The idea of cutting plane methods is to improve the LP relaxation by adding in constraints to cut off non-integer solutions, without cutting off a feasible mixed-integer solution. Solving this improved LP relaxation should then return a solution with no, or less, non-integer values. This can be illustrated by our previous example (3.2). Figure (3.2) depicts the solution space of the problem after adding a constraint $x \geq 3$, which is portrayed by the yellow half-space.

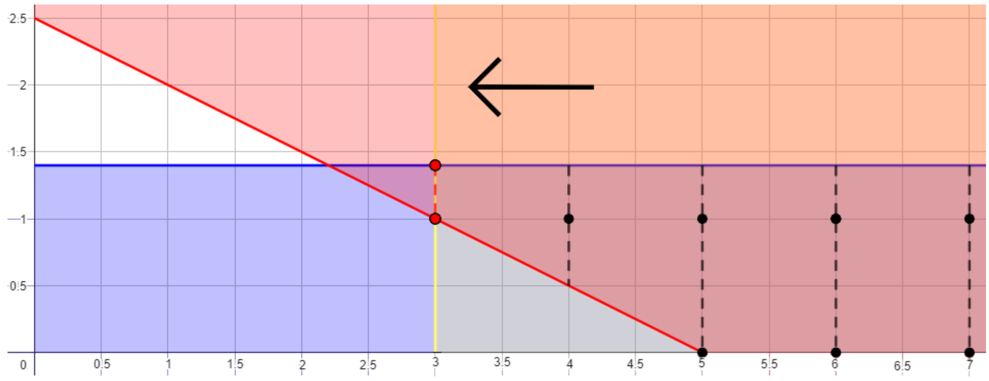

Figure 3.2: Solution space of problem (3.2) with added cutting plane.

Solving the LP relaxation will now return any of the points on the red dashed line (the segment from $(3,1)$ to $(3,1.4)$) as an optimal solution. Any of these points is a feasible solution to the original problem, and hence the cutting plane was effective. In practice, it might not be very easy to find cutting planes that cut off the non-integer solutions that would have otherwise been found. The solution spaces of real problems can be very high-dimensional polyhedrons, and cutting off one infeasible point might very easily cause the LP relaxation to simply find another infeasible point. If instead of $x \geq 3$ we had added any other plane that did not completely cut off everything to left to that same red dashed line (i.e. the set of optimal solutions), we would have still found an infeasible point. Had a plane cut off the entirety of that line, it would not have been valid, as it would cut off all optimal solutions.

Cuts can both be added a priori, or as an iterative procedure. In the latter case, each time the LP relaxation returns an infeasible solution, a cutting plane that cuts off that solution is added. There is not a single method of finding cuts. In fact, Gomory, who originally proposed this method, was amongst those who created different procedures for finding cutting planes. For this, refer for example to chapters 8 and 9 of the textbook Integer Programming by Wolsey [11]. A generic method for finding cutting planes however is the following: take a constraint, or (more likely) combine different constraints by way of addition, subtraction or substitution, such that it contains both integer and non-integer values (variables or parameters). Then, rewrite this so that all the integers are on the left side of the inequality, and the fractional parts are on the right. This right part can now be rounded up or down depending on the direction of the inequality sign, as the left part must be integer. As an example, consider the two first constraints of problem (3.2): $x + 2y \geq 5$ and $y \leq 1.4$. Subtracting the latter twice from the former, and then making a substitution, gives the following inequality: $x \geq 5 - 2y \geq 5 - 2.8 = 2.2$. Rewriting this gives $x \geq 2.2$. Since x is integer, this must mean that $x \geq 3$. This indeed returns the cut we found before, and have seen is valid.

Branch-and-Bound methods

Another method of finding exact solutions is the branch-and-bound method. This method also adds new constraints to the model, but unlike the cutting plane method of cleverly cutting off non-integer solutions outside of the convex hull, this method systematically enumerates, and eliminates, all possible solutions. The broad idea is as follows: if one of the integer-restricted variables of the solution of the LP relaxation is not integer, the value must either be below or equal to that value rounded down, or above or equal to that value rounded up. In case of example (3.2), after the solved LP relaxation resulted in $x = 2.2$, we conclude that either $x \leq 2$ or $x \geq 3$. The branch-and-bound method now creates two new problems, one in which the first constraint is added, and one in which the other is added. Here, the first constraint results in an infeasible problem, and the algorithm concludes that $x \geq 3$ must hold. Since adding this constraint resulted in a fully integer solution, the algorithm terminates.

In practice, this can take much longer, and the method will need to branch multiple times. If we once again suppose that y also needs to be integer, the model can similarly try modifying the problem with the $x \geq 3$ constraint by adding either constraint $y \leq 1$, or constraint $y \geq 2$. An example of a slightly larger application of the branch-and-bound method is given in figure 3.3. The original instance of the problem is P_0. We choose a variable to branch on, and hence create two new problems, P_1 and P_2, that contain an additional constraint that respectively give an upper and a lower bound on this variable. There are four things that can happen. If the solution once more contains a non-integer variable, then we can repeat the procedure and branch again. If the solution becomes infeasible, we prune (i.e. eliminate) this node, because adding more constraints will never result in a feasible solution again. If the solution is feasible and does not contain integer-restricted variables with non-integer values, then two things can happen. Either, this solution has a better objective value than previously found feasible solutions, or its objective value is not as good as that of other feasible solutions found before. If we conclude that a solution cannot be optimal, this node is pruned. It is also possible to prune problems that do still have non-integer integer-restricted variables, if the upper bound (for a maximization problem) of the objective value is worse than the best currently found solution. This is for the same reason that we can prune infeasible nodes: adding more constraints will never allow a sub-problem to have a better objective. To determine if a solution is optimal or not, we need to search all of the other created branches, until all other nodes in the search tree are pruned.

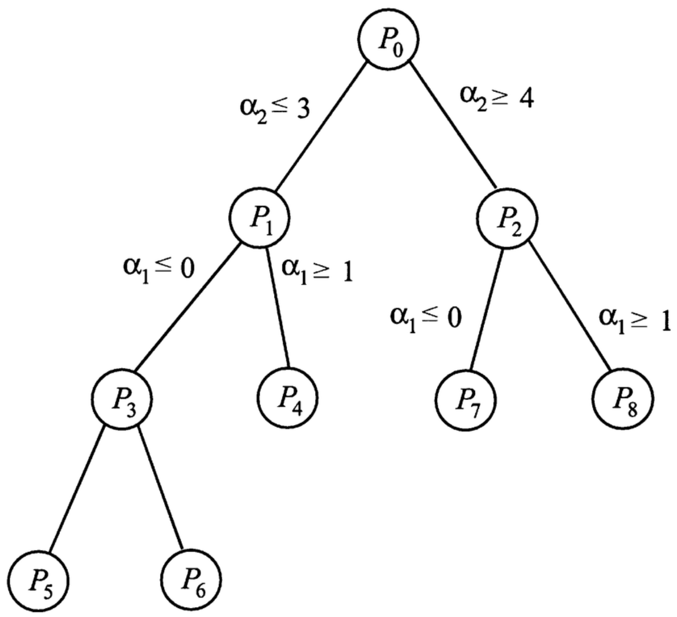

Figure 3.3: An example of a branch-and-bound tree, adapted from Lim and Yu [12] © 2024 IEEE.

This process can take a really long time, as problems may have very many integer-restricted variables, and one can branch on the same variable multiple times. In the worst-case scenario, when no nodes get pruned, the entire search space might get searched. That means we would essentially find the solution by brute-force. For that reason, the performance of the branch-and-bound method cannot be guaranteed. One thing that can impact performance is how the tree gets searched: the sooner a good-quality solution that is feasible for the original (M)ILP is found, the more nodes can be pruned, resulting in a more rapidly shrinking search space. A few possible methods of doing so are depth-first search, breadth-first search and best-first search [13].

Just as important as the searching strategy could be the branching strategy, that determines what variable to branch on. Common branching strategies are most infeasible branching, pseudo cost branching and strong branching. Most infeasible branching simply chooses the variable with the most fractional value (i.e. closest to 0.5), pseudo cost branching keeps track of the variables that have already been branched on to predict which variable will have the greatest likelihood of success, and strong branching tests each of the candidate variables to see which of them gives the most progress [14].

It is also possible to combine cutting plane methods with branch-and-bound, resulting in a method called branch-and-cut. Here, extra cutting planes are added to the subproblems created by the branch-and-bound method. So, instead of only adding a single bound on a variable, more extensive constraints might be added. An example problem that greatly benefits from the branch-and-cut method is the TSP, as its subproblems can be solved quite efficiently. It is also possible to only add cutting planes to the initial LP relaxation, and then branch as usual. This method is instead called cut-and-branch.

Column Generation
Column generation is a technique used to solve large linear programs. Instead of solving the entire LP, only a subset of its variables is considered at first. All the other variables are fixed to 0. Since typically most variables take value 0 anyways (consider for example the TSP, only a small fraction of the possible arcs will be in the final tour), this is not an unrealistic starting point. Then, the algorithm determines what variables are potentially able to improve the value of the objective function. Each variable has a reduced cost value, that says by how much the objective function will improve if the corresponding variable is increased by one unit. In case of a minimization problem, the variable with the lowest reduced cost will bring the greatest improvement of the objective function and will thus will get a non-zero value. The task of finding this variable is often called the (pricing) sub-problem. We will not go into detail on how to solve this, but more information can for example be found in Chapter 11 of Wolsey's textbook on Integer Programming [11]. An example of a problem that benefits from this technique is the set partitioning model formulation of the Vehicle Routing Problem discussed in section 3.2.1. It is also possible to combine this technique with the branch-and-bound method mentioned earlier. This results in a method called branch-and-price, that applies the branch-and-bound technique to the LP relaxation containing only a subset of its variables, and adds variables back in further down the tree when needed.

3.1.4. Metaheuristics
As mentioned before, finding a solution to a (M)ILP is an NP-hard problem. When applied successfully, the exact methods from before do return the optimal solution, but particularly hard problems could easily be intractable. To still obtain a satisfactory solution, if any exist, heuristics must be used instead. Since metaheuristics are more problem-specific, we will not cover these methods in detail. Metaheuristics are typically some variation of a local search technique. Such a technique often succeeds the use of a greedy heuristic that tries to find a good initial solution of the problem. The idea is that you start with an empty solution and each round try to improve some facet of the problem, until the constraints are met and a feasible solution is found. In case of the TSP, a greedy heuristic would add route segments until a completed tour is found. Greedy algorithms can perform well, and are sometimes even optimal, but they are often only used to find a starting solution. Local search heuristics will then try to find improved solutions. To use a local heuristic, one needs to define a neighborhood of solutions that are in some way close to the starting solution. The algorithm then searches this neighborhood for an improved solution. If it can find one, a new neighborhood is defined and the process repeats. If there is none, then the current solution is the found local optimum. In this section, we will not discuss any specific metaheuristics. This is because the models designed for this project are only solved using exact methods. For completeness however, appendix A contains an overview of the most commonly used metaheurisics.

3.2. Mathematical Problems
Now that we have provided sufficient mathematical background, we can go over the Vehicle Routing Problem (VRP) that forms a base for most, if not all, of the vehicle routing algorithms currently in use. From here on, we will look at the variation of the VRP that considers electric vehicles: the Electric Vehicle Routing Problem (EVRP). There also exist many variations of the EVRP, that loosen some of the assumptions posed on the EVRP to describe a more specific and complex problem. We will briefly discuss some of these variations as well.

3.2.1. VRP
Definition
The VRP is a combinatorial optimization problem aiming to calculate an optimal route planning for a fleet of vehicles needing to visit a predetermined set of destinations. Optimal can mean different things for different instances of this problem: the objective function may contain for example the driven distance, the projected duration, or most likely the total cost of a route planning. Just like how there are many variations of the EVRP, the same holds for the VRP. However, since our main interest lies in the EVRP, we will keep the treatment of VRP variations to a minimum.

The first VRP was introduced in 1964 by Clarke and Wright [15], who generalized the "Truck Dispatching Problem" of Dantzig and Ramser five years prior [16] to a linear optimization problem. From that point on, VRP models have gotten much more complicated, trying to incorporate real-life complexities that involve for example time windows for pickup and delivery, or time-dependent travel times [17]. Braekers et al. [17]

describe the classical VRP or Capacitated VRP (CVRP) as follows: the CVRP constructs optimal delivery routes in which each vehicle only travels a single route, where all vehicles have all the same characteristics and emerge from a single depot. The goal is to find a set of routes such that each customer is visited exactly once by a singular vehicle, the vehicles start and finish at the depot and the vehicle capacity is not exceeded, while minimizing the total cost. Table B.1 in appendix B contains an overview of the variations of the VRP that have been studied in the literature. It is outside of the scope of this literature review to thoroughly describe each of these variations, but we do show how the capacitated VRP can be formulated. There are three ways of doing this: using vehicle flow models, commodity flow models, and set partitioning models [18]. Vehicle flow models make use of integer variables that are associated with the set of arcs. When the costs of the solution directly relate to the values of each arc, this method is very convenient. Commodity flow models are similar, but contain an additional set of integer variables that relate to the flow on each arc. They represent the flow of commodities along the routes driven by the vehicles. The set partitioning formulations on the other hand aim to find a minimum-cost collection of circuits that visits each customer once. Note that not all of these formulations have to be modeled as an IP. It is also possible to use dynamic programming [19] for this purpose, as can be seen in the work of Christofides et al. [20].

To illustrate the differences between these methods, we will give three different formulations: a vehicle flow model formulation, a slightly more complicated commodity flow model formulation and a set partitioning model formulation.

Vehicle Flow Model Formulation

There is not one single choice for each of the previously mentioned formulations. For example, for the vehicle flow model one can use two-index variables x_{ij}, that take value 1 if a route visits customer i after customer j and 0 otherwise [21]. Another option is to use three-index variables x_{ijk} that take value 1 if vehicle k visits customer j immediately after customer i [18]. For simplicity, we will give the two-index model formulation presented by Munari et al. [21].

Given a set of customers $\mathcal{C}$, we create a set of nodes $\mathcal{N} = \mathcal{C} \cup \{0, n+1\}$ that include the depot as a departure and a return location 0 and $n+1$. We define the set $\mathcal{A}$ to contain all the (directed) arcs between each of the customers $\mathcal{C}$, and the set $\mathcal{A}' = \mathcal{A} \cup_{j \in \mathcal{N}} (0, j) \cup_{j \in \mathcal{N}} (j, n+1)$ to also include the arcs to and from the depot and the customer. Then, we define the following decision variables:

$$x_{ij} := \begin{cases} 1 & \text{if a route contains arc } (i,j) \in \mathcal{A}'\text{, i.e. visits node } j \in \mathcal{N} \text{ after node } i \neq j \in \mathcal{N}; \\ 0 & \text{otherwise.} \end{cases}$$

$$y_j \in \mathbb{Q}_{\geq 0} : \quad \text{the cumulated demand on the route visiting customer } j \in \mathcal{N}.$$

Let CST_{ij} be the cost associated with traveling on each arc (i, j), K the set of vehicles, DEM_j the demand of each customer j (where $DEM_0 = DEM_{n+1} = 0$), and CAP the maximal capacity of a vehicle. We then get the following formulation:

$$\min_{x} \sum_{(i,j) \in \mathcal{A}'} CST_{ij} x_{ij} \tag{3.3a}$$

s.t.

$$\sum_{j \in \mathcal{C} \setminus \{i\}} x_{ij} = 1 \qquad i \in \mathcal{C}, \tag{3.3b}$$

$$\sum_{\mathcal{N} \setminus \{h, n+1\}} x_{ih} - \sum_{\mathcal{N} \setminus \{0, h\}} x_{hj} = 0 \qquad h \in \mathcal{C}, \tag{3.3c}$$

$$\sum_{j \in \mathcal{C}} x_{0j} = |K|, \tag{3.3d}$$

$$y_j \geq y_i + DEM_j x_{ij} - CAP(1 - x_{ij}) \qquad (i,j) \in \mathcal{A}', \tag{3.3e}$$

$$DEM_i \leq y_i \leq CAP \qquad i \in \mathcal{N}, \tag{3.3f}$$

$$x_{ij} \in \{0, 1\} \qquad (i,j) \in \mathcal{A}'. \tag{3.3g}$$

The objective function (3.3a) minimizes the total cost associated with the chosen routes. The first constraints (3.3b) make sure that all customers are visited exactly once. The next constraints (3.3c) enforce that if a vehicle arrives at a customer h, it must depart from there as well, guaranteeing a correct flow of vehicles. Constraint (3.3d) fixes the number of routes to the number of vehicles available. Together, constraints (3.3e) and (3.3f) make sure that the demand of each customer is met and that the capacity of each vehicle is not exceeded. Since constraints (3.3e) force y_j to increase at every customer (assuming all customers have non-zero demand) this variable will act as an indicator of when each customer has been visited. This will also avoid subtours, i.e. cyclic routes not passing through the depot.

There are different methods of imposing vehicle capacity constraints and subtour elimination constraints. These constraints are in this model formulated by constraints (3.3e) and (3.3f) and were originally introduced by Miller, Tucker and Zemlin [22] for the Traveling Salesman Problem. The advantage of these constraints is that there are only $\mathcal{O}(n^2)$ variables and constraints. This cannot be said about an alternative way of imposing these constraints, one example of which is using the capacity-cut constraints as presented by Toth and Vigo [23]:

$$\sum_{i \in S} \sum_{j \in S} x_{ij} \geq r(S) \qquad \forall S \subseteq \mathcal{C}, S \neq \emptyset. \tag{3.4}$$

The term $r(S)$ denotes the minimal number of vehicles needed to serve all customers in S. These constraints simultaneously impose the connectivity of the solution, and the vehicle capacity requirements. This is a consequence of making sure that a cut $(\mathcal{C} \backslash S, S)$ defined by any set of customers S is crossed by at least as many arcs as $r(S)$. Together with the degree constraints (3.3b) - (3.3d) this implies that each cut $(\mathcal{C} \backslash S, S)$ is crossed just as often in both directions.

The number of constraints of the type (3.4), or any similar set of constraints, grows exponentially with the number of customers. Therefore, the linear programming relaxation of this version of the vehicle flow model can only realistically be solved for very small instances. To partially overcome this it is possible to start only with a limited subset of these constraints, and to add them back in if needed with separation techniques such as branch-and-cut [23]. When we do successfully find a solution to this relaxation, the resulting lower bound does tend to be strong. This immediately shows the main disadvantage of using the constraints we initially presented: the linear programming relaxation might be found easily, but the lower bound it provides is in general significantly weaker than when using the capacity cut constraints (3.4) [24]. We clearly see a trade-off between the complexity of the model and the strength of the linear relaxation [25].

Commodity Flow Model Formulation

The next formulation we give is a commodity flow model formulation. This particular formulation was presented by Toth and Vigo [23], while originally introduced by Garvin et al. [26] and later extended by Gavish and Graves [27], [28]. We will denote our graphs $(\mathcal{C}, \mathcal{A})$ and $(\mathcal{N}, \mathcal{A}')$ in the same way as in the vehicle flow model formulation. We will also choose decision variable x_{ij} the same. We get the following set of decision variables:

$$x_{ij} := \begin{cases} 1 & \text{if a route contains arc } (i,j) \in \mathcal{A}'; \\ 0 & \text{otherwise.} \end{cases}$$

$y_{ij} \in \mathbb{Q}_{\geq 0}$: the first flow variable associated with the arc $(i,j) \in \mathcal{A}'$.

$y_{ji} \in \mathbb{Q}_{\geq 0}$: the second flow variable associated with the arc $(i,j) \in \mathcal{A}'$.

Note that the final two decision variables are closely related. If a vehicle travels from i to j, y_{ij} represents the vehicle load, while y_{ji} represents the vehicle residual capacity along the arc, i.e. $y_{ji} = CAP - y_{ij}$. This is swapped for the arc $(j,i) \in \mathcal{A}'$, so for each arc it holds that $y_{ij} + y_{ji} = CAP$.

We do not need more parameters than we have defined in the previous formulation. However, we will define $TDM := \sum_{j \in \mathcal{C}} DEM_j$ so that we can express the total demand using a single parameter. This results in the following formulation:

$$
\begin{aligned}
\min_{x} \quad & \sum_{(i,j) \in \mathcal{A}'} CST_{ij} x_{ij} && && \text{(3.5a)} \\
\text{s.t.} \quad & && && \\
& \sum_{j \in \mathcal{N}} (y_{ji} - y_{ij}) = 2DEM_i && \forall i \in \mathcal{C}, && \text{(3.5b)} \\
& \sum_{j \in \mathcal{C}} y_{0j} = TDM, && && \text{(3.5c)} \\
& \sum_{j \in \mathcal{C}} y_{j0} = |K|\, CAP - TDM, && && \text{(3.5d)} \\
& \sum_{j \in \mathcal{C}} y_{n+1,j} = |K|\, CAP, && && \text{(3.5e)} \\
& y_{ij} + y_{ji} = CAP x_{ij} && \forall (i,j) \in \mathcal{A}', && \text{(3.5f)} \\
& \sum_{j \in \mathcal{N}} (x_{ij} + x_{ji}) = 2 && \forall i \in \mathcal{C}, && \text{(3.5g)} \\
& y_{ij} \geq 0 && \forall (i,j) \in \mathcal{A}', && \text{(3.5h)} \\
& x_{ij} \in \{0, 1\} && \forall (i,j) \in \mathcal{A}'. && \text{(3.5i)}
\end{aligned}
$$

Our objective function (3.5a) is exactly the same as in the vehicle flow model. Constraints (3.5b) require that the sum over the inflow and the outflow in each customer is equal to twice their demand. The next set of constraints force the commodity leaving the depot to be equal to the total demand (3.5c), the residual load when leaving the depot to be equal to the difference between the total capacity of the vehicles and the total demand (3.5d), and the residual load when entering the depot to be equal to the total capacity of the vehicles (3.5e). These constraints make sure that the commodity flow variables incident to the depot variables behave correctly. Finally, constraints (3.5f) impose that the degree of every customer node is equal to 2 to make sure that each customer can only be visited once.

It has been shown by Baldacci et al. [29] that the linear relaxation of this MILP is stronger than that of the vehicle flow model formulation from before, without the capacity cut constraints. Since we have already seen that the linear relaxation of that formulation is known to be weak, this result is not unexpected [23].

Set Partitioning Model Formulation

The final formulation we are going to present is a set partitioning model formulation. In its simplest form, there is only a single formulation, which was originally proposed by Balinski and Quandt [30] in 1963 and rewritten by for example Munari et al. [21]. In this context, we consider the graph $(\mathcal{C} \cup \{0\}, \mathcal{A} \cup_{j \in \mathcal{N}} (0, j))$ that contains all customers, a single copy of the depot, and all possible arcs in between. We let $\mathcal{H} = \{H_1, \dots, H_q\}$ denote the collection of circuits of this graph starting at the depot and thus representing the entire set of feasible routes, with $q = |\mathcal{H}|$. Associated to each circuit H_j is a cost COC_j, and we have binary parameter VOR_{ij} such that

$$
VOR_{ij} = \begin{cases} 1 & \text{if a vertex } i \text{ lies on route } H_j; \\ 0 & \text{otherwise.} \end{cases}
$$

We then may present the decision variables and the model:

$$
x_j := \begin{cases} 1 & \text{if a circuit } H_j \text{ is selected in the optimal solution;} \\ 0 & \text{otherwise.} \end{cases}
$$

$$
\begin{aligned}
\min_{x} \quad & \sum_{H_j \in \mathcal{H}} COC_j x_j && (3.6a)\\
\text{s.t.} \quad & \\
& \sum_{H_j \in \mathcal{H}} VOR_{ij} x_j = 1 \qquad \forall i \in \mathcal{C}, && (3.6b)\\
& \sum_{H_j \in \mathcal{H}} x_j \leq |K|, && (3.6c)\\
& x_j \in \{0,1\} \quad \forall H_j \in \mathcal{H}. && (3.6d)
\end{aligned}
$$

The objective function (3.6a) minimizes the total cost associated with the chosen routes. The first constraints (3.6b) enforce that every customer gets visited once on one of the chosen circuits. Constraint (3.6c) then makes sure there can only be as many routes as there are vehicles available.

One may note that this formulation is much more succinct than the flow model formulations. This can be explained by the fact that feasibility of a route has been taken care of when designing the set of circuits $\mathcal{H}$ and thus does not require any constraints in the model. The benefit of such a model is that it could be very easy to add further restrictions, such as time windows, to the model as this does not need to be formulated as a constraint. This does lead to the main drawback of this method however: the number of feasible routes can get exponential in terms of the number of customers. Only (very) small instances might be expected to be solved analytically. A column generation approach to solve the linear programming relaxation of the model, followed by a branch-and-price method to find optimal integer solutions, will generally be necessary to find solutions to this model [21]. It should be said that the linear programming relaxation is typically very strong [23], similarly to how we saw that the much larger vehicle flow model formulation also had a significantly smaller integrality gap compared to the initial formulation.

3.2.2. EVRP

Definition

The EVRP extends the VRP by taking battery constraints and charging operations into account. Just like the VRP, it finds a set of vehicle routes starting from and ending at a single depot that visit a set of customer nodes. On top of that however, the routes that vehicles are allowed to take are limited by the battery capacity of the vehicles and the available options for charging. According to Küçükoğlu et al. [5] the basic assumptions for the EVRP can be summarized as follows:

- Each route has to start and end at a depot node.
- Each customer is serviced by exactly one vehicle.
- EVs can visit a charging station to recharge between visiting any two customers.
- Each charging station can be visited by multiple EVs.
- The location of each charging station and their traveling distances from any other node is known.
- The battery level of each vehicle must be kept between 0 and its capacity at all times.
- After visiting a charging station, the battery of a vehicle is fully charged.

Küçükoğlu et al. [5] also mention the most commonly used variations of the EVRP:

- Vehicle capacity restrictions, limiting the weight or volume capacity of a vehicle.
- Time-related restrictions:
  - Time windows for nodes, restricting that a customer must be serviced within a given time window and that each route must be completed within a certain time window.
  - Duration time limits, stating that the total elapsed time for a route cannot surpass the set time limit.
- Partial charging operations, allowing vehicles to only charge partially. Note that this is only relevant when time-related constraints are used.

The following restrictions have also been applied to EVRPs in the past:

- EVRP with Pickup and Delivery
- EVRP with Backhauls
- EVRP with Simultaneously Routing and Siting
- EVRP with Simultaneously Vehicle Recharging and Customer Service
- EVRP with Multiple Depots

The objective function components that are most commonly researched are total number of used vehicles, travel distance and travel time. Other options that have been studied are the total number, or the construction cost, of the used charging stations, total recharging cost or time, total energy consumption, or some other operational costs.

There is also a difference in the the way the energy use is calculated. This can be done with linear deterministic functions based on traveled distance, vehicle speed, vehicle load, road gradient, and more. In the given model, the energy use is calculated as a factor of the distance between two nodes. It is also possible to assign a single energy consumption value to an arc in the graph, or even to combine these methods. Alternatively, stochastic functions can also be considered. Finally, most researchers make use of fleets containing only one vehicle type, but it's also possible to use heterogeneous fleets that contain different types of EVs, and possibly even ICE vehicles. In section 4.1.2, this will be treated in more detail.

Basic EVRP

We start by giving a formulation for the most basic model, and then consider how to adapt this for the most common different variations. This model has been taken from the work of Küçükoğlu et al. [5], which was derived from several other papers on variations of the EVRP. We do need to introduce some additional notation. For convenience, we give a complete list of the required sets, parameters and decision variables:

- $0, N+1$: Depot nodes
- F: Set of charging stations
- F': Set of dummy nodes to allow for multiple visits to the nodes in F
- $\mathcal{C}$: Set of customers $\{1, \ldots, N\}$
- $\mathcal{C}_0 := V \cup \{0\}$: Set of customers and depot node 0
- $\mathcal{C}_{N+1} := \mathcal{C} \cup \{N+1\}$: Set of customers and depot node $N+1$
- $\mathcal{C}' := \mathcal{C} \cup F'$: Set of customers and charging stations
- $\mathcal{C}'_0 := \mathcal{C}' \cup \{0\}$: Set of customers, charging stations and depot node 0
- $\mathcal{C}'_{N+1} := \mathcal{C}' \cup \{N+1\}$: Set of customers, charging stations and depot node $N+1$
- $\mathcal{C}'_{0,N+1} := \mathcal{C}' \cup \{0\} \cup \{N+1\}$: Set of customers, charging stations and depot nodes $0, N+1$
- K: Set of vehicles
- LEN_{ij}: Length, or traveling distance from node i to node j, $\forall i, j \in \mathcal{C}'_{0,N+1}$
- ECR: Energy consumption rate of the vehicles per unit of distance
- VBC: Battery capacity of the vehicles

The decision variables and model formulation are then the following:

$$x_{ij}^k := \begin{cases} 1 & \text{if vehicle } k \text{ travels from node } i \text{ to node } j,\ l_{ij} > 0; \\ 0 & \text{otherwise.} \end{cases}$$

$$y_i^k \in \mathbb{Q}_{\geq 0} : \quad \text{battery level of vehicle } k \in K \text{ on arriving at node } i.$$

$$
\begin{aligned}
&\min_{x} \quad \sum_{i\in\mathcal{C}'_0}\sum_{j\in\mathcal{C}'_{N+1}}\sum_{k\in K} LEN_{ij}x_{ij}^k && && (3.7a)\\
&\text{s.t.} && && \\
&\sum_{j\in\mathcal{C}'_{N+1}}\sum_{k\in K} x_{ij}^k = 1 && \forall i \in \mathcal{C}, && (3.7b)\\
&\sum_{j\in\mathcal{C}'_{N+1}}\sum_{k\in K} x_{ij}^k \le 1 && \forall i \in F', && (3.7c)\\
&\sum_{j\in\mathcal{C}'} x_{0j}^k \le 1 && \forall k \in K, && (3.7d)\\
&\sum_{i\in\mathcal{C}'_0} x_{ij}^k = \sum_{i\in\mathcal{C}'_{N+1}} x_{ji}^k && \forall j \in \mathcal{C}', \forall k \in K, && (3.7e)\\
&y_j^k \le y_i^k - (ECR\cdot LEN_{ij})x_{ij}^k + VBC(1-x_{ij}^k) && \forall i \in \mathcal{C}, \forall j \in \mathcal{C}'_{N+1}, \forall k \in K, && (3.7f)\\
&y_j^k \le VBC - (ECR\cdot LEN_{ij})x_{ij}^k && \forall i \in F'\cup\{0\}, \forall j \in \mathcal{C}'_{N+1}, \forall k \in K, && (3.7g)\\
&y_0^k \le VBC && \forall k \in K, && (3.7h)\\
&x_{ij}^k \in \{0,1\} && \forall i,j \in \mathcal{C}'_{0,N+1} (i\neq j), \forall k \in K, && (3.7i)\\
&y_j^k \ge 0 && \forall j \in \mathcal{C}'_0, \forall k \in K. && (3.7j)
\end{aligned}
$$

The objective function aims to minimize the total traveled distance. Constraints (3.7b), (3.7d), (3.7e) have the same function as constraints (3.3b), (3.3c) and (3.3d) respectively. The difference is that the current model uses three-index decision variables, the extra index of which is used to indicate the vehicles. Constraints (3.7b) make sure that all customers are visited exactly once by one of the vehicles. Constraints (3.7d) limit the number of routes to the number of vehicles, and constraints (3.7e) guarantee a correct flow of vehicles. The vehicle flow model of the VRP did not contain charging stations, so there are no analogous constraints to constraints (3.7c) that ensure that each dummy charging station can be visited only once. However, constraints (3.3e) and (3.3f) do resemble constraints (3.7f) and (3.7g). In the vehicle flow model formulation these constraints took care of the customer demand, while in the basic EVRP they model the battery demand of the routes. Additionally, constraints (3.7h) make sure that the battery of a vehicle is fully charged after leaving the depot.

Again based on the work by Küçükoğlu et al. [5], we describe how to alter this model to obtain the EVRP with partial charging, the capacitated EVRP, and the EVRP with time windows.

EVRP with a Partial Charge Policy
To allow vehicles to charge only partially, we need to introduce a new decision variable:

$q_i \in \mathbb{Q}_{\ge 0}$: battery level of the vehicle before departing node $i \in F' \cup \{0\}$.

We may then introduce constraints (3.8) and (3.9)

$$y_j^k \le q_i - (ECR\cdot LEN_{ij})x_{ij}^k + VBC(1-x_{ij}^k) \quad \forall i \in F'\cup\{0\}, \forall j \in \mathcal{C}'_{N+1}, \forall k \in K \tag{3.8}$$

$$y_i^k \le q_i \le VBC \quad \forall i \in F'\cup\{0\} \tag{3.9}$$

to replace constraints (3.7g) and (3.7h). The consequence is that at each charging station (or the depot) the vehicle only gets charged up to value q_i, which is fixed to be between the battery level upon arrival of the charging station, and the maximal battery capacity.

Capacitated EVRP (CEVRP)
If we want to restrict the weight or load capacity of the vehicles, we need to introduce a new parameter. Just like in the VRP, let CAP be the weight/load capacity of the vehicle. We already defined DEM_i to be the demand of customer $i \in \mathcal{C}$. Constraints (3.10)

$$\sum_{i \in \mathcal{C}} \sum_{j \in \mathcal{C}'_{N+1}} DEM_i x_{ij}^k \leq CAP \quad \forall k \in K \tag{3.10}$$

make sure that the total capacity of each vehicle is not exceeded. These can simply be added to the given formulation to gain the CEVRP.

EVRP with Time Windows (EVRPTW)
Adding time windows to the formulation is a slightly more complicated affair than the previous alterations. For that purpose, we define new parameters:

- TRT_{ij}: Travel time from node i to node j $\forall i, j \in \mathcal{C}'_{0,N+1}$
- EST_i: Earliest time to start service at node $i \in \mathcal{C}'_{0,N+1}$
- LST_i: Latest time to start service at node $i \in \mathcal{C}'_{0,N+1}$
- SET_i: Service time at node $i \in \mathcal{C}_0$
- RCR: Recharging rate of the batteries

and a new decision variable

$$s_i \in \mathbb{Q}_{\geq 0}: \quad \text{start time of service at node } i \in \mathcal{C}'_{0,N+1}.$$

We can then add the following constraint to the model:

$$s_i + (TRT_{ij} + SET_i) \sum_{k \in K} x_{ij}^k \leq s_j + LST_0(1 - \sum_{k \in K} x_{ij}^k) \quad \forall i \in \mathcal{C}_0, \forall j \in \mathcal{C}'_{N+1}; \tag{3.11}$$

$$s_i + TRT_{ij} x_{ij}^k + RCR(VBC - y_i^k) \leq s_j + (LST_0 + RCR \cdot VBC)(1 - x_{ij}^k) \forall i \in F', \forall j \in \mathcal{C}'_{N+1}, \forall k \in K; \tag{3.12}$$

$$EST_i \leq s_i \leq LST_i \quad \forall i \in \mathcal{C}'_{0,N+1}. \tag{3.13}$$

These constraints track the duration of all operations, determine charging times and make sure the resulting time windows are feasible, when either using a full or partial charging policy. Constraints (3.11) ensure that the time windows when leaving customers are complied with, and the same holds for constraints (3.12) for charging stations. If we allow partial charging, we use the variable q_i to determine how long a battery is charged for, instead of assuming it needs to be charged until full. Then, we should use constraints (3.14) instead of constraints (3.12):

$$s_i + TRT_{ij} x_{ij}^k + RCR(q_i - y_i^k) \leq s_j + (LST_0 + RCR \cdot VBC)(1 - x_{ij}^k) \forall i \in F', \forall j \in \mathcal{C}'_{N+1}, \forall k \in K. \tag{3.14}$$

The last constraints (3.13) enforce the time windows of each of the nodes.

3.3. Solution Methods

As we have seen in section 3.1, there is not a single way to solve an Integer Program. For simpler or smaller problems, using exact methods might be a good and accurate method of finding solutions. When a problem becomes more complicated due to (very) large input data or added complexity, the running time could become too big to find a solution in a practical amount of time. Then, it might no longer be the desired method. Since Integer Programming is NP-hard, there is no guarantee that a solution will be found in any reasonable amount of time. Instead, metaheuristics that try to approximate a solution are used for such cases. This section will touch upon both of these solution methods to see how researchers solved their version of the EVRP.

3.3.1. Commercial Solvers

There is a small subsection of papers that to find solutions of their model depend on the use of commercial solvers. The most popular of these are CPLEX and Gurobi. Basso et al. [31], Soysal et al. [32] and Kopfer and Vornhusen [33] use CPLEX as their solution method, while Chen et al. [34] and Wang et al. [35] do the same with Gurobi. Other researchers solve small instances of their model with commercial solvers, while large instances can be solved with proposed metaheuristics. This is done for example by Xiao et al. [36] and Lu et al. [37] with CPLEX, and Kancharla [38] and Froger et al. [39] with Gurobi.

3.3.2. Branch-and-Bound Approaches

Branch-and-bound and its variations are exact methods, and therefore are still only suitable for less complicated problems. Nonetheless, there are a few papers that use such a method as a main solution method. Lee et al. [40] use the branch-and-price approach, and Tahami et al. [41] use branch-and-cut. More commonly these methods are combined, resulting in the so-called branch-price-and-cut method. This is applied by for example Ceselli et al. [42], Munari et al. [43] and Costa et al. [44].

3.3.3. Heuristic Approaches

To solve larger and more complicated problems, exact solution method will no longer be sufficient. There is a broad variety of methods one can use to solve variations of the EVRP, such as applying some of the standard metaheuristic solution methods described in appendix A. As we do not discuss these methods in the literature review itself, a discussion about the application of these methods by different authors can be found in appendix C. That chapter additionally includes some examples of hybrid metaheuristic approaches, and also mentions some alternative approaches that have been used to solve these types of problems in the past.

4 Battery Modeling

This chapter discusses the different ways in which the battery of an electric vehicle has been modeled in the literature. The first section considers the energy consumption of the vehicle battery while driving, while the second section covers the opposite behavior: charging. We will describe how these components can be modeled, and how important specific parts of the model are.

4.1. Energy Consumption

In this section, we discuss different ways of modeling the energy consumption of batteries, and how this is impacted by different factors. As the battery is a critical element of the EVRP, in section 3.2.2 we already saw a simple method of modeling the battery charging and discharging. Then, we used constraints (3.7f) - (3.7h) to model the battery capacity in the same way customer demand was modeled before using constraints (3.3e) and (3.3f), where we defined the energy demand to simply be the distance between two locations multiplied by some energy consumption parameter. While these constraints do approximate the battery level, there are many ways of improving this. Vehicle-specific factors such as the shape of the car and how much extra weight it contains, as well as the behavior of the driver (e.g. at what speed is driven, how often and quickly the vehicle brakes, etc.) and external factors such as the road incline and temperature, all have a direct impact on how quickly a battery discharges. A single parameter value will never be able to take all these complexities into account. Many research papers do approximate energy consumption linearly [5], either as a function of distance, time, or vehicle load. Depending on the number of parameters and their sophistication, linear approximations of energy consumption can get more or less realistic.

4.1.1. General Model

The starting point for modeling the energy consumption of an EV is generally determining the mechanical power. The way this is calculated is often credited to Barth et al. [45], that determined the engine power demand of diesel trucks. This calculation is broken down neatly by Goeke and Schneider [46]. The mechanical power is the power P_M in W needed to overcome the rolling resistance, aerodynamic resistance and gravitational force. In order to express these forces, the following values are needed:

- m: Total vehicle mass in kg
- g: Gravitational constant in m/s^2
- c_r: Rolling friction coefficient (depends on tire pressure and road surface conditions amongst other factors)
- θ: Gradient angle in degrees
- ν: Velocity in m/s
- c_d: Aerodynamic drag coefficient
- ρ_a: Air density in kg/m^3
- A_f: Frontal area of the vehicle in m^2
- a: Acceleration in m/s^2

The rolling resistance F_r (in N) can be determined as

$$F_r = c_r \cdot m \cdot g \cdot \cos(\theta). \tag{4.1}$$

The aerodynamic drag F_a (in N) is determined by

$$F_a = \frac{1}{2} \cdot \rho_a \cdot A_f \cdot c_d \cdot \nu^2. \tag{4.2}$$

Finally, the gravitational force F_g (in N) is given by

$$F_g = m \cdot g \cdot \sin(\theta). \tag{4.3}$$

In addition to the sum of these forces that add up to the traction force F_T, Newton's second law $F_A = m \cdot a$ is also included in the calculation. To go from a sum of forces to the power, these terms are multiplied with the velocity ν. That results in the following total power demand:

$$P_M = \big(m \cdot a + \frac{1}{2} \cdot \rho_a \cdot A_f \cdot c_d \cdot \nu^2 + m \cdot g \cdot \sin(\theta) + c_r \cdot m \cdot g \cdot \cos(\theta)\big) \cdot \nu. \tag{4.4}$$

As done by Barth. et al. [45], this power requirement still needs to be turned into demanded engine power requirement. This transformation is done by dividing P_M by the energy efficiency of transmission, motor and power conversion ε_M and adding the engine power demand caused by running losses of the engine and using vehicle accessories like air conditioning P_{acc}. That gives total power demand P_{out} in W:

$$P_{out} = \frac{P_M}{\varepsilon_M} + P_{acc}. \tag{4.5}$$

An ability that diesel trucks do not have is regenerative braking. When driving downhill or decelerating, the traction force F_T may become negative and some of the energy can be transmitted back to the battery. Asamer et al. [47] describe how to model this recuperated energy. In order to be able to recuperate any lost energy, the vehicle does need to be driving at a certain minimal speed ν_{min}. Sandrini et al. [48] take ν_{min} to be $15\ km/h$. As vehicles usually drive at much higher speeds, this limitation is not given much attention in the studied literature on the topic. For even a moderately simplified theoretical model, it however is an important consideration. We define ε_G as the efficiency of transmission, generator and in-vehicle charger. With that, we can describe the regenerated energy after braking:

$$P_{in} = \begin{cases} 0 & \text{if } \nu \leq \nu_{min}; \\ F_T \cdot \nu \cdot \varepsilon_G + P_{acc}. & \text{otherwise} \end{cases} \tag{4.6}$$

For vehicles that have a single motor powering two wheels, if the braking strength is below $0.2g$, all braking force gets allocated to the wheels connected to the electric motor. When braking harder, this strength is allocated to all four wheels, and some of the energy is lost due to friction brakes [49]. However, as the majority of deceleration is below that limit [47], this is an appropriate measure nonetheless.

Not all papers take regenerative braking into account. In part this is for simplicity, but it can also be a deliberate choice for a different reason. An example of this is the work by Pelletier et al. [50], which presents a robust MILP for the EVRP. Not accounting for an extra inflow of energy is a simple tool to make a problem more robust, as this allows for a little more leeway for the driver when the predicted battery level gets low.

In order to find the total energy demand E in J of a trip, we need to integrate the total power P over the duration of the trip. At any moment in time,

$$P = \begin{cases} P_{out} & \text{if } F_T \geq 0; \\ P_{in} & \text{if } F_T < 0. \end{cases} \tag{4.7}$$

If we define the duration of the trip by T and express the time with variable t, we may conclude that the total energy demand is the following:

$$E = \int_0^T P\ dt. \tag{4.8}$$

By multiplying this result with $3.6 \cdot 10^6$, the unit of E becomes kWh. If we denote the total power demand at the i'th timestep by P_i, where i is an integer ranging from 0 to $\lceil \frac{T}{\Delta t} \rceil$ for a fixed stepsize Δt, then this integral can be approximated by the function $E = \sum_i P_i \cdot \Delta t$. It is also possible to calculate the energy consumption per unit of distance by dividing the total energy consumption of a trip by the distance of that trip. The parameter ECR as defined in section 3.2.2 that described the energy consumption rate can thus be derived from the general formula $\frac{E}{L}$, in which L is the length of a single trip.

Seeing how the energy consumption is expressed as a discretized integral, the main term of which can vary greatly between only small moments of time, we need a way of breaking this down further. In 2019, Basso et al. [51] presented a method of calculating the energy expenditure between road links between two intersections. Each road link is split up in three parts, and the energy consumption is calculated for each of them: an acceleration phase ($e^{\uparrow}$), a steady speed phase ($e^{\rightarrow}$) and a braking phase ($e^{\downarrow}$). In figure 4.1, these phases are illustrated. Note that a similar thing can be done if the start and/or end velocity is not equal to zero. Hulagu and Celikoglu [52] illustrated this in more detail. Another paper that uses such an approach to take acceleration and deceleration into account is the work of Pelletier et al. [50], although they do not describe it nearly as thoroughly.

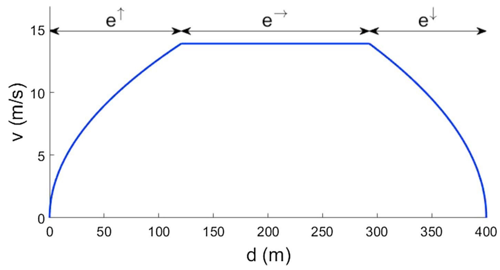

Figure 4.1: Speed curve from standing still to a complete stop with steady speed 50 km/h. Source: [51].

However the energy consumption of a vehicle is expressed, it is immediately clear that it cannot be substituted directly into a vehicle routing model. Even with endless computational power, there is no way a model will be able to account for the exact driving behavior of each driver in every situation. Therefore, the need for making approximations arises. To learn what needs to be approximated, the next section describes the way energy consumption in EVs has been modeled in the past. Section 4.1.3 will then discuss methods to find suitable parameters.

4.1.2. Practical Models

This section will discuss how different papers studying different variations of the EVRP modeled energy consumption previously. This can broadly be split into three categories: linear deterministic functions, nonlinear deterministic functions and stochastic functions.

Linear Deterministic Functions

Using linear deterministic functions to approximate the energy consumption of an electric vehicle can be done in a great many ways. We divide this in four different categories, ranging from very simple to quite refined.

Single Energy Consumption Parameter We saw that the simplest way to model this was by using a single energy consumption parameter that predicts the energy consumption solely based on the distance. Papers that have used this method are for example written by Keskin and Çatay [53], Roberti and Wen [54], and Chen et al. [34]. A computationally equivalent approach is basing the energy consumption not on the distance, but on the trip duration. This method however isn't very popular: Küçükoğlu et al. [5] only identified 4 papers out of 136 publications that used this method. The most recent one of which was published by Yang et al. [55]. A common factor amongst these papers is that the research interest is purely computational, and realistic battery modeling is not the priority.

Energy Consumption Parameter per Vehicle If a model uses a mixed fleet, some of the modeled coefficients may differ between the different vehicles. Therefore, it might not be reasonable to assume that each vehicle spends the same amount of energy on each unit of distance. A simple way of extending the previous method is by assigning different energy consumption parameters to different vehicles. This is done for example by Hiermann et al. [56]. If the amount of time that a vehicle takes to charge is not (or only slightly) constrained, a similar effect can be achieved by keeping the energy consumption rate the same, but assigning different battery capacities to different vehicles, as done by Zhau and Lu [57].

Energy Consumption Parameter per Arc A slightly more involved, but still computationally acceptable approach is calculating the energy consumption parameter for every arc in the graph. This can significantly boost performance if different trajectories differ greatly in terms of road type and incline. This can be applied to homogeneous fleets, an example of which is the work of Barco et al. [58] that studies an airport shuttle service scenario. It is also possible to consider a heterogeneous fleet, as done by Li et al. [59].

Energy Consumption as a Function of the Vehicle Load More commonly, papers that opt for a more realistic battery consumption model use a slightly more complicated but potentially much more accurate approximation by taking the vehicle load into account. The difference in weight between an empty and a fully loaded vehicle can be massive and might significantly impact energy expenditure. To implement this linearly, the battery discharge will need to be modeled as a linear function of the vehicle load at a given timestep. This is done by for example Futalef et al. [60] and Pelletier et al. [50] for a single vehicle type. Kopfer and Vornhusen [33], and Goeke and Schneider [46] similarly used a fleet of heterogeneous vehicles.

Nonlinear and Stochastic Functions
Energy consumption can be approximated fairly well with linear functions, since many important factors do not vary between solutions. A consequence is however that the values of the decision variables (that can vary between solutions) can only have limited impact on the energy consumption. For the modeling variations discussed thus far this did not pose a problem, but when trying to incorporate the effect of the state of charge (SoC) on the energy consumption rate, a linear approximation is not sufficient. Kim and Chung [61] noticed that the SoC of the battery did have a measurable difference on the energy expenditure, and proposed a model in which the energy consumption (in addition to the charging process) is a nonlinear function of the SoC.

While most researchers assume the energy consumption to be deterministic, there are a few papers that treat this stochastically. The reasoning for this is clear: accurately estimating the energy consumption is a complicated and detailed task depending on a large multitude of factors that are impossible to take into account exactly. To combat this, the uncertainty of the estimation is built into the model, that can in turn optimize the problem in such a way that a solution is sufficient for the worst case scenario. A common method for doing this is making a model robust, such as Pelletier et al. [50] did. They achieve this by using the deterministic energy consumption estimation as the expected value, and adding a randomly distributed uncertainty term. The probabilistic part of this term is derived from a carefully chosen uncertainty set, designed in such a way that aims to reflect reality best. Soysal et al. [32] use a similar approach, except that they make the assumption that the battery consumption is distributed normally.

A slightly different method to achieve the same goal was used by Basso et al. [62] in 2021. They published a paper building on their previously designed model [51], adapting it to a stochastic model using full Bayesian regression techniques. Just like the previous two methods, they were able to use a MILP solver to find routing solutions. Another method expresses the uncertainty in energy consumption by using fuzzy numbers: Zhang et al. [63] defines a fuzzy optimization model, based on uncertainty theory. The resulting model however is not a MILP, so instead they propose their own algorithm.

4.1.3. Parameter Estimation
In the previous section, the different ways in which energy consumption was expressed in an EVRP were discussed. In order to successfully implement this, it is necessary to use parameters that accurately describe the situation. This section will examine how this can be done. Example values for the parameters that are commonly not highly variable are given in table 4.1. Note that the values for some of these parameters depend on the vehicle. As an example, the values related to the medium $75kWh$ edition of the Citroën ë-Jumpy will be used.

Table 4.1: Typically elementary model parameters.

Total vehicle mass (excluding driver) m	$1,988$ - $3,105$ kg
Gravitational constant g	9.81 m/s^2
Rolling friction coefficient c_r	0.02
Aerodynamic drag coefficient c_d	0.4
Air density ρ_a	1.127 - 1.164 kg/m^3
Frontal area of the vehicle A_f	3.81 m^2

Less Variable Parameters The total vehicle mass of the electric truck is between $1,988kg$ when empty, and $3,015kg$ when at max capacity, as presented by the manufacturer [64]. An estimate of the rolling friction coefficient for car tires on asphalt is given by the Engineering Toolbox website [65]. Chowdhury et al. [66] estimate the aerodynamic drag coefficient c_d for this type of delivery truck to be around 0.4. The Engineering Toolbox website gives values for the air density ρ_a for a wide range of temperatures. The values in the table are for temperatures of $30°$ C to $-10°$ C. The frontal area of the electric truck follows from the dimensions of the truck given by the manufacturer [64]. The gradient angle θ depends entirely on the driven road, so a single general parameter cannot be given. It is however information that can be known in advance, and therefore not considered to be highly variable.

More Variable Parameters Other parameters depend on both the vehicle and factors such as temperature, and are less trivial to estimate. These include the efficiencies of the different power conversions in the vehicle ε_M and ε_G. The power demand of vehicle accessories P_{acc} can also be treated as a parameter, and can be harder to estimate than the previous two terms, as this depends very much on the driver also. Asamer et al. [47] solve this issue by considering an interval between 236 W and a maximum of 1266 W, choosing 450 W as the default value. Fiori et al. [67] assumed a value of 700 W, and also considered situations with higher auxiliary power demands when the temperature outside is either very cold or very warm. Basso et al. [51] also considered different options, one of which was 20 kW, much larger than the other estimates. This is unrealistic for short trips in a regular vehicle, but it does approximate the auxiliary power demand of large trucks with contents that need to be kept frozen.

Often, ε_M and ε_G are taken as constants. Barco et al. [58] and Yuan et al. [68] fix these values to 0.9, Wu et al. [69] use 0.95, Travesset-Baro et al. [70] use a much lower value of 0.69. Asamer et al. [47] respectively consider values between 0.63 - 0.9 and 0.64 - 0.82 within their sensitivity analysis, but set a baseline value of 0.9 and 0.8 respectively. Basso et al. [51] did not pick a single value, but used regression analysis to find a function of the velocity describing the efficiency. These methods are however lacking, as the range of a vehicle can differ greatly depending on temperature: differences of 36% of the median have been observed during winter [71]. An approach that avoids this problem is to simply ignore these factors and use big data to model the energy consumption rate directly, as done by Fetene et al. [72].

We can clearly see that there are many ways of choosing these parameters, and that there is not a single approach that results in success. The same thing holds for the velocity and acceleration of the vehicle. Unlike the previous terms, these values correlate strongly with the time. Whereas the weight of the vehicle may differ only between stops, the speed and acceleration may differ anywhere on the route and can depend on unpredictable factors. The end of section 4.1.1 described ways of breaking down the energy consumption into smaller parts based on the acceleration and deceleration of the vehicles. When using a speed profile, the velocity and acceleration are assumed known at all times, and the energy consumption can be calculated. Hulagu and Celikoglu [52] created their own speed profiles, while Basso et al. [51] referred to existing tools to simulate a driving cycle. Another option is to use data from driven vehicles, like Fetene et al. [72] did, which can be very detailed and hence result in accurate outcomes. The price of this accuracy however is flexibility; the resulting models will only be completely valid for the vehicles and roads the data is sourced from.

Specific Effects

Up until now, this section explained how EV batteries are generally modeled. This is a very challenging task, especially when the information available is limited. When estimating the energy consumption of vehicles, it can be very helpful to get an idea of how different relevant factors impact this process. This section will discuss how differing conditions or modeling assumptions affect the battery consumption of an EV.

Effect of Temperature and Weather Conditions The outside temperature can have an enormous effect on the performance of the battery. Independent papers found a difference of 36% [71] and 34% [72] of the predicted energy consumption between summer and winter. Fetene et al. [72] also calculated the correlation between the energy consumption and the temperature, wind and precipitation. The temperature is negatively correlated with the energy consumption, while the wind and precipitation have positive correlation. Remarkable is that even when correcting for these factors, there is a positive correlation between winter and energy consumption.

The sensitivity analysis performed by Asamer et al. [47] concluded that variance in the motor efficiency ε_M can also have a significant impact on the performance of the model. Seeing how this efficiency is determined in part by temperature, this is no surprise. The experiments of Xie et al. [73] concertize this statement. They compared the energy consumption under ambient temperatures $(-20°C, -10°C, 0°C, 25°C, 45°C)$. With baseline temperature $25°C$, they found that colder temperatures can increase the energy consumption by respectively 58.02%, 40.40% and 33.31%. At $45°C$, the energy consumption decreased by 1.93%. Iora and Triboli [74] found similar results: compared to an ambient temperature of $20°C$, a temperature of $0°C$ and $-15°C$, the range of the EV dropped with respectively 40% and 60%.

Donkers et al. [75] studied the impact of ambient temperature on energy consumption for different speeds, and concluded that this effect mainly holds for lower velocities. When driving at $130\ km/h$, the relative influence of the temperature is significantly lower. They also concluded that for higher temperatures, the vehicles become more energy efficient. This confirms the results of Xie et al., although the negative correlation between temperature and energy efficiency is not found for driving at a much lower speed of $30\ km/h$. A related result of Donkers et al. is that when the vehicles have to counteract strong headwind speed, the energy efficiency plummets, in particular for high speeds: when driving at $130\ km/h$ with a headwind speed of $100\ km/h$, about three times as much energy is used compared to a windless scenario. For lower speeds a similar factor is found, but due to the difference in energy consumption at different speeds, the effect is stronger for high velocities.

Effect of Elevation Asamer et al. [47] also considered effects for different elevation profiles. Noteworthy conclusions are that only when driving on a negative slope, does varying the recuperation efficiency ε_G make a significant difference in the outcome. Variance of the total mass of a vehicle results in negligible differences when driving on a flat or negatively sloped road, but when the road becomes hilly or positively sloped, different values can make a large difference. Liu et al. [76] directly studied the impact of the road gradient on the performance of the battery consumption model. They concluded that when elevation is considered in the model, the resulting outcome is $5-8\%$ more accurate than without. Donkers et al. [75] also studied the impact of the slope on the energy consumption, and found that a slope of $1°$ compared to a completely flat road increased the energy consumption by about a third for vehicles driving at $130\ km/h$. This agrees with the simulation created by Genikomsakis and Mitrentsis [77]. When driving at $30\ km/h$, the effect was much stronger: Donkers et al. found that the energy consumption almost doubled. Al-Wreikat et al. [78] did a similar study, and found that ascending roads with a slope of $1.72°$ increased the specific energy consumption by 50%, while a descending road with that same slope, the specific energy consumption decreased by 80% compared to a flat road. While this these differences seem large, practical behavior can be different than these experiments. A case study performed in the mountainous region Andorra by Travesset-Baro et al. [70] estimated that when accelerating and braking moderately, a 21.8% drop in EV range is expected. With an aggressive driving style, this increases further to a 26.9% reduction.

Effect of Acceleration and Speed The link between speed, acceleration and energy consumption is directly clear from the way that they are calculated, as the velocity is the only variable that appears as a high-order term in the function for the total power demand (4.4). Consequently, driving at a high velocity causes a disproportionate jump in energy usage. This phenomenon however is not explicitly researched by much of the literature, since most researchers use predetermined driving profiles to base their experiment on. Results varied between different cycles. One study that specifically investigated different driving cycles was for example written by De Gennaro et al. [79]. To gain insight in the relationship between the velocity and the efficiency, we refer to papers that also use simulations to draw conclusions. It should be noted that the acceleration and speed of a vehicle is a large descriptor of the driving style of an individual. The next section will shortly discuss driving styles, and their effect on the electrical efficiency of the vehicle.

Basso et al. [51] considered the differences between using a detailed speed profile, versus taking the average speed. They concluded that the difference in energy consumption was very large, in particular at higher speeds (say, above $50\ km/h$) and shorter driving intervals (accelerating and breaking withing less than $1\ km$). They expressed the relationship between the energy consumption estimated at an average speed (e_{avg}), and the energy consumed when accelerating and braking are considered too (e), by their ratio $\frac{e_{avg}}{e}$. Figure 4.2 calculates this ratio as a function of the driven length, at multiple different speeds. For example, when driving 600 meter with on average $60\ km/h$, there is roughly a factor 5 difference. Thus, it is insufficient to simplify the model in such a way, especially when a driver needs to accelerate and brake frequently on a route.

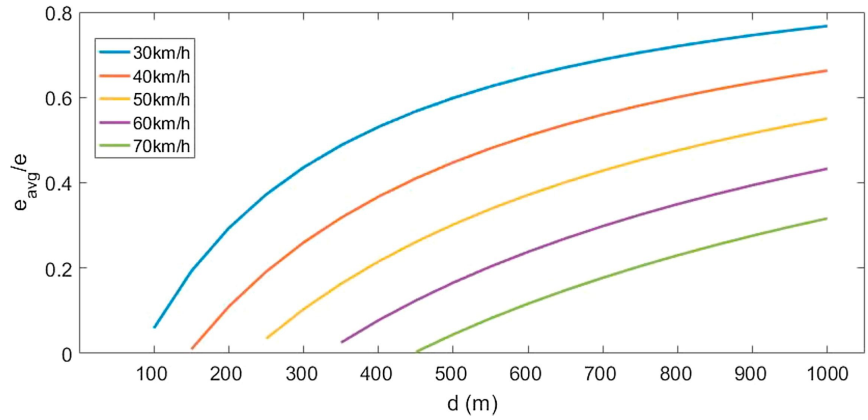

Figure 4.2: The ratio between the average energy use and the energy use over the entire cycle e_{avg}/e for different vehicle speeds and trip lengths $d(m)$. Source: [51].

Mamarikas et al. [80] simulated the energy consumption as a function of the average speed. Up to about $40\ km/h$, the energy consumption in Wh/km rapidly decreases. From that point on however, it slowly increases again. When vehicles are driving $140\ km/h$, the energy consumption is predicted to be roughly equal to vehicles driving on average around $5\ km/h$ (due to frequently stopping and re-accelerating). Conventional ICE vehicles generally drive at optimal efficiency at a much higher speed than $40\ km/h$. Note that if regenerative braking is kept out of the calculations, the simulation implies that driving at $80\ km/h$ is almost exactly as efficient as driving at $40\ km/h$, therefore in part explaining that difference. Another simulation confirms that the recovered energy from regenerative braking is largest around roughly $20-45$ km/h. This is in line with our comment about equation (4.6) about ν_{min}, which was given an example value of $15\ km/h$.

Effects of Other Factors Fetene et al. [72] concluded that long trips ($\geq 10\ km$) were more battery efficient than short trips ($\leq 2\ km$): on average short trips consume $40\ Wh/km$ more than medium trips and $57\ Wh/km$ more than long trips. Duarte et al. [81] studied the effect of the battery SoC on energy use of a full hybrid EV, and concluded that with lower SoC levels, the energy consumption was higher. These outcomes illustrate the weakness of modeling the energy consumption as a function of (only) the distance. Hulagu and Celikoglu [52] show that there is indeed a difference between a solution that aimed to find the shortest paths, and a solution that aimed to be as energy efficient as possible.

Fetene et al. [72] also studied the difference in energy consumption between driving on highway roads, versus driving on non-highway roads. Correcting for the different driving behaviors in both situations, they found that this characterization was statistically insignificant. El Amrani et al. [82] on the other hand found that EVs are more efficient on trunk roads than on highways. The sensitivity analysis of Asamer et al. [47] concluded that the rolling friction coefficient does have a high influence on the model outcome. The rolling friction coefficient depends on a lot of different factors, such as the tire, the tire pressure, the velocity, road surface conditions, etc. Even when ignoring the difference amongst roads (to a reasonable degree), it is still important to select this parameter carefully. Jonas et al. [83] also studied the different energy consumption efficiences at different roads, and concluded that a smart choice of road can potentially save up to 46% of energy usage. However, they noticed the highest energy consumption on between interstate roads and local roads, as these respectively allowed for high speeds and frequent traffic interruptions. Unlike Fetene et al. however they did not correct for these differences in the driving profile.

Doyle and Muneer [84] modeled the energy consumption of both heating and cooling systems in EVs, both during the daytime and nighttime. For heating, they found values for the energy consumption of around $5\ Wh/minute$ when the difference in temperature was around $10°C - 15°C$, up to $55\ Wh/minute$ for a temperature difference of $24°C$. As a result, the heating accounted for up to 30% of the energy consumption. On average, this came down to 18%. In the case of cooling, they found that when the cabin was cooled at least $2°C$, a minimum of $3\ Wh/minute$ was used. This number went up to almost $30\ Wh/minute$ when the the desired temperature was $9°C$ cooler than the outside temperature. On average the cooling resulted in a 14% share of the total energy consumption. Fiori et al. [67] did a similar study, but assumed a range for acceptable temperatures $(17°C - 24°C)$ instead of heating the vehicle to the upper end of this range, and cooling it to the lower end as in the previous work. Heating a vehicle up from $-5°C$ resulted in a 10% to 32% increase in energy consumption, which agrees with the previous results. When cooling the vehicle from $25°C$ and $35°C$, the energy consumption increased with respectively 1% to 3% and 3% to 11%. This difference can be explained entirely that the previous results assumed the vehicle to cool to $17°C$, and therefore needed much more cooling.

Another factor that can be accounted for is driving style. Donker et al. [75] identified three driving styles: eco-driving, normal driving and aggressive driving. These categorizations were made based on the use of regenerative brakes, the speed, acceleration and deceleration. At high speed ($130\ km/h$), the researchers found a 17% increase in energy consumption for the aggressive driver, compared to the eco-driver. At speeds below $30\ km/h$ on the other hand, the eco-drivers used up to 5% more energy, which can be explained by the higher energy consumption of the heating and cooling system in the vehicles.

One paper that explicitly researched the impact of different loads on the efficiency is the work of Mruzek et al. [85]. They considered loads ranging from $50\ kg$ to $250\ kg$, with steps of $50\ kg$, and found that with the heaviest load, the energy consumption was 7% higher than with the lightest load. Moreover, the Depth of Discharge (DoD) measured for each of these loads behaved linearly. This means that the assumption of an energy consumption factor as a product of the added weight used in the energy consumption prediction models appears to be realistic.

The final factor we consider in this section is that of congestion, for which Mamarikas et al. [80] created a model. At low speeds, their simulation concluded that moving from a congested to a normal scenario reduces energy consumption by 15%, while the other way around causes an increase of 6%. At high speeds, we respectively see an increase of 2% and a decrease of 20%. Measures to reduce congestion, such as mini roundabouts or other traffic calming methods, were also found to reduce the energy consumption by 2% up to 28%.

4.2. Charging

This section will cover the different available charging techniques, and how these variations result in different model formulations.

4.2.1. Charging Techniques

When calculating the energy consumption of an EV, the technology inside the battery is, at least with the currently available battery technology, not crucial. For the charging behavior of an EV it could be critical what charging system is used, in part due to the wide range of available options. Most commonly, vehicle batteries use conductive (wired) charging. A logistically similar charging solution is using static inductive (wireless) charging. To avoid the waiting time coinciding with in particular low-voltage charging, another possible technique is using battery swapping. A method with even greater practical benefit is using dynamic inductive charging. This comes down to making roads that contain coils under the asphalt, transferring energy directly to the vehicle allowing it to charge while driving.

Dynamic Inductive Charging and Battery Swapping

Most of these methods are still very futuristic. There have been a few experiments of dynamic inductive charging performed by Electreon [86], in which they electrified small stretches of roads up to $2\ km$ long.

The technology is immensely promising, but still under development, not to mention the huge investments it requires. Until its implementation takes up speed, it does not make sense to consider this as a main method of charging. It's not even certain that this technology is feasible to apply broadly. If that does happen, the model formulation will start to look different, as vehicles might arrive at a location with more battery than when they left the previous location. The battery swapping technique has similar issues. The main difference is that the implementation of this technique is further along: NIO, one of China's largest automakers, was operating as of May 2023 over 1360 battery swapping stations in China, and 13 in Europe [87]. Despite that, there are still many practical hurdles to overcome before this method of charging can become mainstream. Nevertheless, the EVRP with battery swapping has a moderate base of research. To give a few examples: Chen et al. [34] and Raeesi and Zagrafos [88] studied a variation of this problem with time windows. Soysal et al. [32] instead considered a pickup and delivery problem that used battery swaps. Li et al. [89] does consider the standard EVRP, and focuses on the energy consumption and carbon emissions. Finally, Meng and Ma [90] considered both battery swapping and traditional charging stations in their research. The benefit of using battery swapping in a model is that it closely relates to the full charging policy that is a part of the most basic EVRP variation as described in section 3.2.2. Due to a lack of availability, we will not consider these methods of charging.

Static Inductive Charging and Conductive Charging
While static inductive charging is just like its dynamic counterpart still in development, there already are companies offering inductive charging pads. A few big players are InductEV [91], WAVE Charging [92], WiTricity [93] and Plugless Power [94]. Almost no EVs allow for for inductive charging out-of-the-box, but it is possible to install systems that do make this possible. This means that as of yet, this technique, even if available, is not accessible for the average EV driver. Static inductive charging can offer great benefits, not only in convenience but due to a reduction in charging time as well. Unlike inductive phone chargers that are known to lose a lot of energy due to heat, inductive EV chargers [95] can reach similar (or even better!) efficiencies than conductive chargers [96]. It is therefore likely that the number of static inductive chargers will grow over the coming years. Luckily, they can be modeled similarly to conductive chargers, likely up to (some of) the parameter values, as both methods require the driver to visit a particular place and wait. For that reason, this research will focus entirely on the most widespread method of charging: conductive charging.

Conductive Charging Levels
Conductive charging can be based on multiple different technologies. Generally however it is split into three categories: level 1, level 2 and level 3 charging [97]. Level 1 uses $120V$ AC charging, which makes it suitable for household outlets in North America. It draws 1.4–$1.9\ kW$ power, making it the slowest method of charging, needing roughly 8 to 16 hours to fully charge a small vehicle. European homes typically have $230V$ power supplies, and can thus make use of single-phase level 2 charging. This is achievable for Americans too, by upgrading to a $240V$ power outlet. This charging level can provide 7.7–$25.6\ kW$ power and charge a small EV in 4 to 8 hours. Level 2 charging also includes three-phase $400V$ AC charging, which is suitable only for public installation. Level 3 charging is better known as DC fast charging. Level 1 and 2 convert AC power to DC using an on-board charger, only then allowing it to enter the battery. DC fast charging omits this step: The power is converted from AC to DC off-board, charging the battery directly. As this bypasses the on-board charger and the limitations that come with it, the battery can be charged much more quickly. In fact, it's possible to fully charge an EV in 15 minutes. A charging time of 30 minutes is however more realistic, since most vehicles cannot charge at the highest power: DC fast charging can range from $50\ kW$ to $350\ kW$ [98]. As this is achieved by using voltages in the range of $400V$ - $1000V$, this is only suitable for commercial or specific public applications.

Clearly, the speed at which a vehicle can be charged depends greatly on the available charging infrastructure, in addition to the capabilities of the vehicles themselves. The differences between levels however go further than simply speed. The charging curve, describing how fast a battery is able to charge at different levels of SoC, differs significantly between AC and DC charging, as illustrated in figure 4.3. Thus, to accurately model the charging, it is crucial to know what kind of charger is used.

Figure 4.3: Charging curve for AC and DC charging. Source: [99].

4.3. General Model

The type of chargers available are not the only variable when determining how to model the charging in an EVRP. Restricting to conductive charging, Küçükoğlu et al. [5] identified four different policies: using full, or partial charging, in addition to allowing a single, or multiple different charger types. They found 22 papers using the most complicated method: partial charging and allowing different charger types. Most papers however limited themselves to allowing vehicles to only use a single type of charger.

Montoya et al. [100] were the first to introduce non-linear charging functions to the EVRP in 2017. As they were the first to deviate from the (then) standard way of modeling charging, they had to argue how to do it instead. This section is based in part on their reasoning.

A typical charging scheme used for rechargeable batteries is a constant current-constant voltage (CC-CV) charging scheme [101]. As the name suggests, during the first (CC) phase the charging current is kept constant, causing the SoC to increase linearly over time until the second phase. The CV phase starts after the terminal voltage of the battery reaches a specified maximum, generally around 80% SoC. From this point on, allowing the voltage to increase could cause serious battery degradation, due to the rise in temperature amongst other factors [102]. Therefore, the terminal voltage is held constant, which means that the current will exponentially decrease, as will the increase in SoC. This is illustrated in figure 4.4. CC-CV is not the only existing charging scheme; Hemavathi and Shinisha [103] describe a total of 8 charging schemes, exhibiting a large variety in possibilities. Since CC-CV is used often and the increase of the SoC is moderately easy to approximate, we will only focus on this method.

While the shape appears to be easy to approximate, is it very complicated to model analytically. Not only factors such as current and voltage determine the charging behavior, so do self-recovery [105] and temperature [100]. As a result, the battery SoC is described using differential equations. These however are very difficult to integrate into an EVRP model, instigating the need for approximations. Wu et al. [106] for example describe the charging behavior analytically, but discussing these models is outside the scope of this literature review. For completeness however, we do give a simplified time-discretized model, as provided by Pelletier et al. [101]. Their model states that the battery level SOC_k at a certain timestep k can be determined as follows:

$$SOC_{k+1} = SOC_k - \frac{\Delta t}{3600VBC} \cdot i_k. \tag{4.9}$$

Here, Δt is the length of the timestep, VBC is the capacity of the battery in Ah, and i_k is the current in A. During the CC phase, i_k is constant, resulting in a linear function. For the CV phase however, this does not hold. Instead, we calculate i_k as follows:

$$i_k = \frac{V_{OC}(SOC_k) - V_{CV}}{R}. \tag{4.10}$$

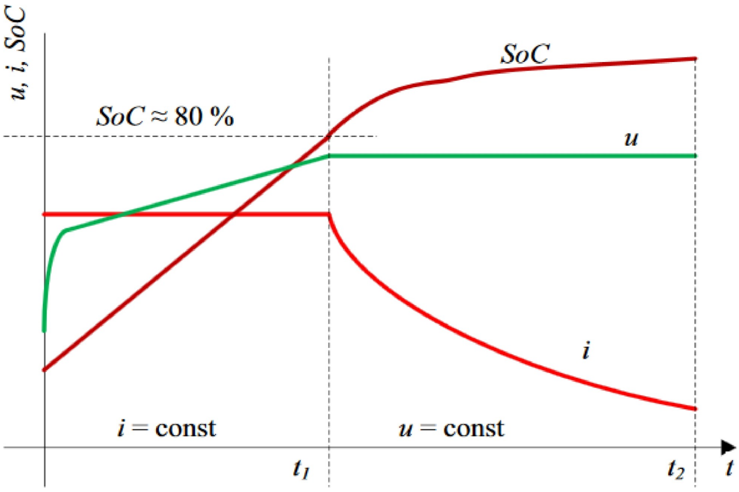

Figure 4.4: Depiction of the CC-CV charging scheme. The x-axis depicts the time t, such that the CC phase ends at t_1, while the CV phase ends at t_2. Here, i, u respectively represents the current and the terminal voltage. Source: [104] © 2024 IEEE.

$V_{OC}(SOC_k)$ is the open voltage of the battery, given as a function of the SoC. One way of expressing this is as follows: $V_{OC}(SOC_k) = a_1 e^{-a_2 SOC_k} + a_3 + a_4 SOC_k + a_5 e^{\frac{-a_6}{1-SOC_k}}$, of which parameters $a_1, \ldots, a_6$ were experimentally determined and are given by Pelletier et al. [101]. V_{CV} is the terminal voltage at the moment the CV phase starts, and R is the internal battery resistance. According to the results by Wu et al. [106], this value does not only depend on temperature, it also varies slightly based on the SoC. Clearly, even this simplified model is not suitable for use within an EVRP.

Common Approximations

Montoya et al. [100] identified three methods of linearizing the charging function. The first of these is called first-segment approximation. As the name suggests, in this approximation only the first segment, the CC phase of the charging scheme, is considered. This is for example done by Bruglieri et al. [107] in their research. When only charging up to 80%, this approximation is fairly accurate. More commonly however, the entire charging duration is linearized. This is done for example by Keskin and Çatay [53]. Typically, the calculations for the approximations are not given. Montoya et al. considered two options. Either, the charging rate corresponds to the slope of the linear segment of the charging function, or the charging rate is the slope calculated between the first and last measurement. This first approximation (L1) is too optimistic, as in the final charging phase the charging rate is highly overestimated. The final approximation (L2) has the opposite problem, as for the majority of the charging process it highly underestimates the real charging rate. All of these approximations are pictured in figure 4.5.

Both of these methods have considerable drawbacks: either the vehicle charges for longer than necessary, negatively affecting the quality of the solution, or the vehicle does not charge long enough. In the worst case scenario this results in having a driver be unable to finish a route due to an overly optimistic battery charge estimate. To improve these approximations, Montoya et al. [100] decided to instead approximate the charging curve $g_i(y_i, \Delta_i)$ with a piecewise-linear function. y_i denotes just like in the basic EVRP model in section 3.2.2 the battery level arriving at node i, and Δ_i stands for the time spent charging at this node. Before doing so however, they rewrite this function to contain only the single index m denoting the number of timesteps a vehicle is charged for. Now, $g_i(y_i, \Delta_i) = \hat{g}_i(\Delta_i + \hat{g}_i^{-1}(y_i))$. Based on the data of Uhrig et al. [108], they approximated $\hat{g}(m)$ with three pieces, for charging speeds $11kW$, $22kW$ and $44kW$. These approximations respectively had average relative absolute error 0.90%, 1.24%, and 1.90%, indicating high accuracy. Figure 4.6 shows their approximation for charging a $16kW$ battery with a $22kW$ charger.

At the start of this section we argued that using a single approximation for both AC as well as DC charging was insufficient. In most of the literature however, this distinction is not made. There could be a few reasons for this. For one, AC charging is still the norm, and most often available. Secondly, not all vehicles, especially those manufactured a few years ago, are able to use the fastest (DC) charging

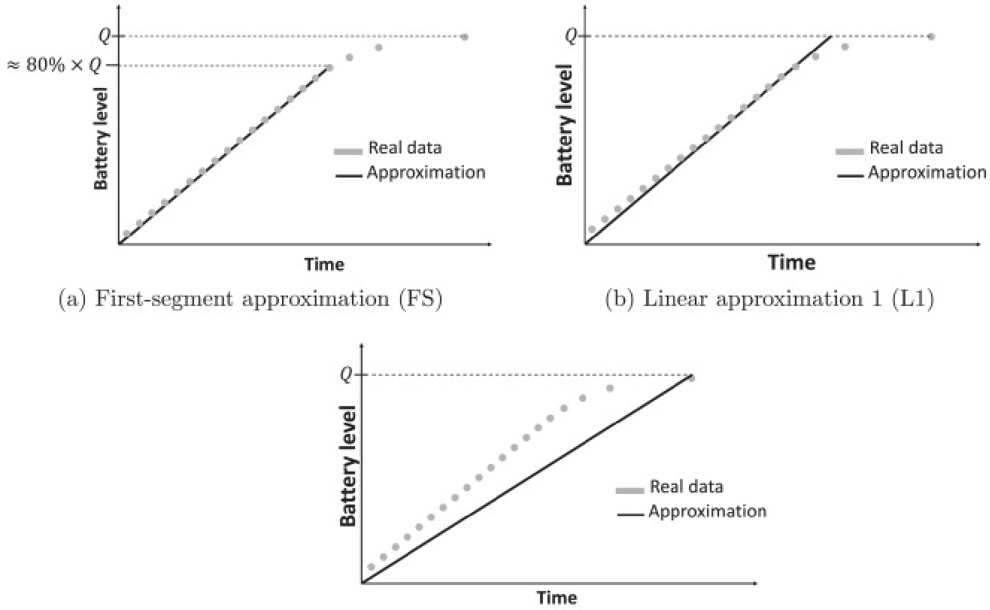

(a) First-segment approximation (FS)

(b) Linear approximation 1 (L1)

(c) Linear approximation 2 (L2)

Figure 4.5: Linear approximations in the literature compared to real data. Source: [100].

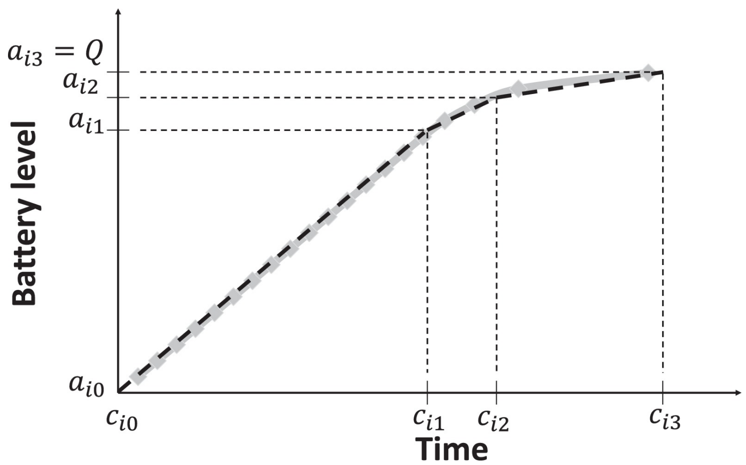

Figure 4.6: Real data versus piecewise linear approximation of a $22kW$ charger charging a battery of $16kWh$. Source: [100].

methods. Fast charging is also known to accelerate battery degradation [101], although this was likely not a main decision factor. What could have played a role was that AC charging is much easier to approximate (piecewise) linearly. Modeling AC charging behavior can already be tricky due to the changes in power near the end of the charging cycle, but it still behaves mostly linearly. In figure 4.3 it can be seen that DC charging behaves very differently. The constantly changing power output implies that on no part of the SoC domain, a constant charging rate will provide an accurate estimate. If the same technique as shown in figure 4.6 would be used, a much more complicated approximation would be needed. For that reason, a completely different approach is likely necessary when modeling DC fast charging in the EVRP. As DC fast charging is not a critical aspect of our research, we will leave the question on how to model this open.

4.3.1. Practical Models

This section will discuss how the different methods of calculating the charging behavior that were mentioned in the previous section were applied in variations of the EVRP. The simplest option is to only allow vehicles to charge overnight at the depot. When we do allow mid-route charging, we already saw that we can split the modeling choices into four categories based on using a full or partial charging policy, and using only a single charging speed or allowing for multiple. We also saw that while many papers approximate the charging behavior linearly, it can be done much more accurately using a piecewise linear function.

No Intermediate Charging
There are two ways of dealing with the limited range of an EV. Either, you make sure that the vehicle returns before it runs out of battery, or you make sure to schedule a stop somewhere that allows the vehicle to increase its range. The former option comes at the cost of being limited to shorter routes, but it does reduce both the cost of electricity as well as the time worked by the driver. Xiao et al. [36] additionally expect the range of vehicles to increase by so much that charging during the day is not necessary, therefore choosing this approach. Erdoğdu and Karabulu [109] instead simplified the model in this way to be able to study the relationship between their two objectives: minimizing the total distance, and the total energy consumption. Abdallah et al. [110] also studied a version of the EVRP without charging stations, instead focusing on the relationship between the range of a vehicle and the driven speed.

Full Charging Policy, Single Charging Method
The next simplest way of modeling charging is by demanding each vehicle leaves a charging station fully charged, assuming homogeneous charging speeds. This can once more be split into two different methods. Either a fixed amount of time is reserved for charging every single time, such as done by Lu et al. [37] and Li et al. [111], or a fixed rate is used to determine how long it will take before the vehicle is fully charged, such as Booth and Beck [112] did. Kopfer and Vornhusen [33] used the former method, although they realized that the final 20% of the charging cycle takes disproportionally long, and decided to only charge up to 80% on recharging stations.

Full Charging Policy, Multiple Charging Methods
In this variant of modeling the charging, the vehicles once more leave each charging station with a full battery. This time however, there isn't a single fixed charging rate or time. Küçükoğlu et al. [113] consider two different charging rates, allowing both quick and very slow charging. Basso et al. [51] do not specify the rates, but do allow each charging station to have a unique charging rate. Conrad and Figliozzi [114] take a slightly different approach. Instead of having distinct charging stations, they define a subset of the customers to function as charging stations. They also defined a single parameter γ that determines when a vehicle is 'fully' charged. For $\gamma = 0.8$, each vehicle leaves the customer with 80% battery, which takes 0.8 times the complete charging time that was determined for that location.

Partial Charging Policy, Single Charging Method
Using a partial charge policy, even when fixing a single charging rate, already allows for a lot more freedom in the model. This method can be seen as the simplest option of the more advanced modeling choices. Keskin and Çatay [53] apply this by solving the EVRPTW with a single charging rate. Goeke and Schneider [46] also aim to solve the EVRPTW but now consider a mixed fleet of electric and conventional vehicles. A similar problem is studied by Li et al. [59], that also considered simultaneous pick-up and delivery service. A final example is the work by Yang et al. [55], studying the EVRP with mixed backhauls. While their model formulation allows for a unique charging rate per charging station, they only experiment with varying the charging rate for all stations at the same time.

Partial Charging Policy, Multiple Charging Methods
We can add slightly more detail to a problem when allowing for multiple charging methods. One way of interpreting this is combining for example partial charging with battery swapping, which is exactly what Mao et al. [115] did. As our focus however lies on conductive charging, and we assumed static inductive charging could be modeled in the same way, we consider multiple charging methods to mean different charging rates. Ceselli et al. [42] assumed that every charging station has a subset of different charging technologies available, each with a different recharging speed and a recharging unit cost. Keskin and Çatay used a similar setup, except that they assume that every charging station has every technology available. One of these technologies in fact is level 3 charging ($50 - 100kW$), although just like the slower methods they approximate the charging process linearly. An example of a slightly more complicated variation of this version is the work by Chen et al. [116]. They do assume that the charging behaves linearly with an inversely correlated cost function, but the whole problem is formulated as a stochastic optimization problem with chance constraints. Similarly, Yao et al. [117] give vehicles the same freedom of charging, but instead aim to model the EVRP together with monetary incentives.

Non-Linear Charging
Every single paper that we've encountered that models charging in a non-linear way, references Montoya et al. [100]. It may also be noted that there isn't a single paper that uses a non-linear approximation for the charging behavior, and chooses to model a full-charging policy. Once more can we differentiate between using a single charging method, or allowing multiple options.

Single Charging Method Futalef et al. [60] used the method of Montoya et al. [100] to solve their EVRP with capacitated charging stations with a genetic algorithm. Unlike Montoya et al. however, they only used a single charging function, instead of multiple, for different speeds. Karakatič [118] used the same approach to instead solve the multi-depot EVRPTW. Lee et al. [40] also uses a piecewise linear approximation when modeling the charging when solving the EVRP, but this work set itself apart because unlike Montoya et al. it aims to find solutions exactly, without using any form of heuristic. A side effect of this newfound method is that instead of only allowing piecewise linear approximation functions, this method will use any concave and non-decreasing function as a charging function.

Multiple Charging Methods Froger et al. [39] studied the same problem as Montoya et al. [100], the EVRP with nonlinear charging. They introduced two new formulations that they showed are more effective than the existing formulations. Kancharla and Ramadurai [38] extended the EVRP with nonlinear charging from Montoya et al. by including load-dependent discharge. This different formulation, together with enhanced solution methods, demonstrated better performance. Koç et al. [119] also extended the formulation of Montoya et al., this time by considering shared charging stations. This variant assumes that several companies jointly invest in charging stations that could offer charging at different speeds, and tasks the model at finding optimal locations.

4.3.2. Battery Degradation
A discussion about modeling battery charging is not complete without mentioning battery degradation. Different papers treat this topic differently. Some papers, in particular papers that strongly simplify the charging aspect of the model, hardly, or not at all acknowledge this issue. One example is the work by Li et al. [111] that assumes a constant charging time for every single charging operation. Other papers do take the battery lifespan into consideration, albeit indirectly. Futalef et al. [60] tries to preserve the lifespan of the modeled batteries by adhering to a so-called State of Health (SOH) policy. This policy consists of constraining the SoC of the battery between fixed values, whenever possible. Here, they chose the lower and upper bound to respectively be 38% and 82%. This does significantly impact the range of the EV, as less than half the battery is allowed to be used. Seeing how models usually occupy the full (or a larger interval of the) battery capacity, it seems like this price is too much to pay for many, as a severe reduction of an already less-than-ideal range can be insufficient to perform the needed operations. Some papers however go further than this, and add a battery degradation cost to the objective function, such as Barco et al. [58] did. They modeled this cost as a sum of the degradation due to temperature, SoC and DoD, the details of which are outside of the scope of this literature review.

4.3.3. Parameter Estimation and External Effects
Unlike the energy consumption, the difficulty with appropriately modeling the charging behavior does not rely on finding and applying the right information. Instead, the arduous part of modeling the charging is finding the actual formulation and particularly how to efficiently solve that formulation. Nonetheless, the two processes do have in common that researchers that are mainly interested in the modeling, and less in the applicability, simply choose a seemingly suitable parameter for both charging as well as discharging. For more specific problems, it is likely known what charging methods are available. It should be known how quickly a specific charger is able to charge a vehicle, and what the highest speed is that an individual vehicle can be charged at. More specifically, the vehicle owner should know what technologies the vehicle is compatible with. If this information is not known, there is no other information that could help find a better approximation.

There are two ways in which this conclusion is insufficiently nuanced. Firstly, in section 4.3 a crude battery charging model was provided, as modeling the entire charging cycle is very difficult to do analytically. To properly approximate the final part of the charging cycle, more information is necessary. Montoya et al. [100] and subsequent related papers made their piecewise linear approximation based on data from Uhrig et al. [108]. Froger et al. [39], Karakatiç [118], Kancharla et al. [38] and Koç et al. [119] appear to indeed have mimicked this approximation exactly. Futalef et al. [60] on the other hand did not specify how the parameters for their piecewise linear charging function were derived. Lee et al. [40] instead of a piecewise linear function used a logarithmic function to simplify the charging behavior, derived in a similar way as was done in section 4.3. This clearly requires some additional parameters, but as this method is to the best of our knowledge applied in only a single paper, we will not elaborate.

Secondly, while the charging process does not vary as much as energy consumption does, there are still external factors that can determine how quickly a vehicle charges. In section 4.3 we saw for example that the charging behavior depends on the internal battery resistance, which in turn depends on the temperature and even to some degree on the SoC. Lindgren and Lund [120] studied the effects of different ambient temperatures on the charging efficiency, expressed as the self-weighted mean charging power (SWMCP). According to them, the SWMCP is a more faithful expression than defining the charging capability in terms of energy stored per unit time. They found that the SWMCP is 15% lower at $-10°C$ compared to $20°C$. Figure 4.7 contains the values of the SWMCP for the entire range of researched temperatures, both without standby options, and with battery thermal management in use. The difference between these two circumstances is not very large, but it can be noticed that the variability is generally larger in the first case. Based on these results, the modeling accuracy could be improved if in (very) cold scenarios, the determined charging rate decreases by up to 15%.

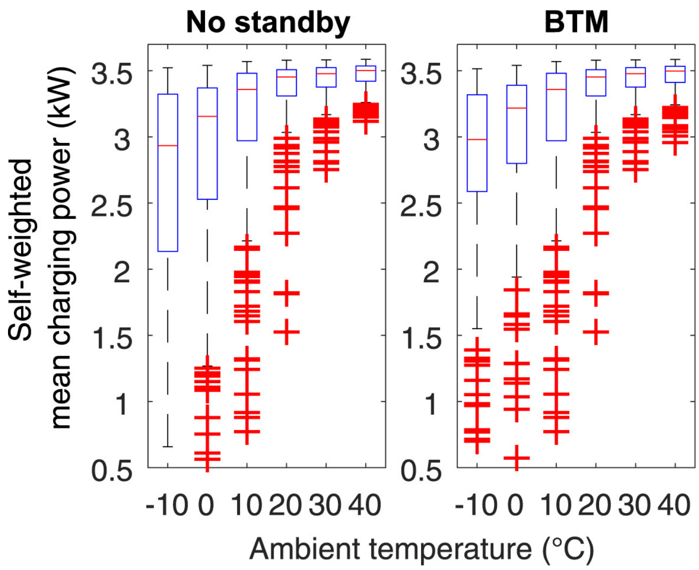

Figure 4.7: Self-weighted mean charging power for $3.6kW$ charging at different ambient temperatures with and without Battery Thermal Management, adapted from Lindgren and Lund [120].

It should be noted that it is not certain that these results can be extrapolated to faster charging methods. Similar research has been done by Motoaki et al. [121], who instead studied the effects of cold temperatures on $50kW$ DC fast charging. They found an even stronger effect for this type of charging. With 95% confidence, a 30-minute charging session at $0°C$ was expected to result in a $22-36\%$ decrease in SoC compared to that same session at $25°C$. When a vehicle is charged with this technology, compensating for a drop in temperature is thus even more important than for lower-speed AC charging.

5

Contribution to the Literature

This final section will quickly summarize the contents of this literature review, and after that will explain how this work aims to contribute to the existing literature.

5.1. Conclusion

In this literature review, we discussed the most commonly used methods of modeling different variations of the vehicle routing problem and modeling both the charging and discharging behavior of an electric vehicle. Chapter 3 defined Mixed-Integer Linear Programming, a broad method used to express and solve all kinds of optimization problems, one of which is vehicle routing. We mentioned different solution methods, in particular ways of solving these kinds of problems exactly, and glossed over the use of metaheuristics that might not find an optimal solution, but could perform very well nonetheless. In the next section we introduce the VRP, and its very broad group of variations. For this problem, we described multiple formulations and already noticed that solving instances exactly might not be feasible. Continuing onto the EVRP, we saw how the formulation of an instance of the VRP can be made suitable for electric vehicles, by adding constraints that mimic the energy consumption and charging of the vehicle battery. We also saw how this formulation can be adapted to a few different variations. Finally, this chapter explained different methods that are used to solve these problems in the literature.

Chapter 4 contains two main sections, a part on modeling the energy consumption of a battery, and a part on modeling the battery charging process. The first section first explains the general model that is used to model energy consumption of a vehicle, specifically an electric vehicle. It then describes a few different methods that EVRP formulations use to model the energy consumption. This ranges from using a single energy consumption parameter, to using a nonlinear function of both vehicle-specific as arc-specific properties. Next we discuss the parameters that need to be determined in order to calculate the energy consumption according to the given model. Some of these are very easy to determine, such as the gravitational constant or the frontal area of the vehicle, but the power efficiency parameters and an applicable speed profile can be challenging. Finally, the effect of varying some of these parameters or related factors is considered. There are many factors that can significantly impact the resulting prediction. The most important factors to predict carefully are the temperature, the behavior of the driver (when the driver speeds up and brakes, and how much heating or cooling is used) and the elevation (in particular when driving through mountainous regions).

The next section discusses the methods that are used by researchers to model the charging process of an EV. The section starts with an overview of the different charging techniques that are currently available or in development. Then, we give a coarse model describing the charging method that is most commonly used. As charging a battery is a highly non-linear process, even this simplified model is unsuitable for implementation into a MILP. For that reason, this behavior is very often linearized. We describe three methods that can be used for this purpose. Furthermore, we also describe a commonly used piecewise linear approximation that can model the charging process with high accuracy. Next, multiple methods of incorporating charging into the EVRP are provided. The simplest method is assuming vehicles will only charge at the depot, discarding the need for charging stations, while the most complicated method allows non-linear approximations for charging at different speeds. We quickly discuss battery degradation, and finally comment on the used parameters and the effect of temperature on the charging process.

5.2. Contributions of this Research

In this thesis, the goal is to create an algorithm that can calculate an optimal routing and charging schedule, under the assumption that each vehicle will drive two (or more) shifts on the same day (that may consist of multiple trips) and can only charge in between those shifts at a depot. We additionally assume that the charging capacity is limited. To keep the settings realistic, we typically assume to have access to a set of chargers larger than strictly necessary. Only occasionally are the limits of the available chargers fully tested. This algorithm will be a combination of two separate algorithms: A scheduling algorithm that assigns vehicles to both chargers and shifts, and a routing algorithm that creates routes based on the energy levels of the vehicles.

To the best of our knowledge, such a variant of the EVRP has not been studied before. In the current literature on electric vehicle routing, the charging needs outside of those during the driven routes are usually not considered in this manner. There certainly are papers discussing EVRP variants with capacitated charging stations in either of the ways described above. However, as these charging stations are generally not located at the depot, they are much more spread out and will only need to facilitate a subset of the vehicles. The nature of our variant is thus very different, as the charging possibilities are much more restricted and therefore require more sophisticated scheduling than without those restrictions.

Additionally, this research applies to both industry as academic interests. One of ORTEC's clients wants to increase the number of electric vehicles in their fleet. Charging the vehicles during the day is one way for them to use these vehicles in a more efficient manner. The models designed in this project and the conclusions drawn from the results can therefore inform this company, and other companies in similar situations, about the effects and practicalities of charging vehicles during the day in between shifts. This problem is of academic interest as in this form, it is a new variation of the EVRP, introducing new techniques and ideas about how EVs can be used in practical situations. In the discussion section, a few extensions to the existing models are suggested. These extensions have little practical use, as their goal is to find improved solutions when the set of available chargers is very limited. In reality, no company would have installed less chargers than they expect to need. Trying to solve such complicated instances can however be a very interesting mathematical challenge.

Part II

Model Design

6 Preliminary Work

This chapter contains everything that is needed to design our model. We will first discuss what technique will be used to implement this model. The next thing to discuss is what data we have access to, and how it applies to our problem. Finally, we will determine what our main objectives are. As we are not collaborating directly with any clients, we can determine these ourselves.

6.1. Starting Model

To experiment with different options of modeling both the charging scheduling as well as the routing of EVs, we need an engine that can solve both scheduling and routing problems. Initially, we considered three different options. The first option was to use AIMMS [122], an advanced development environment for building optimization based operations research applications, to design problems solving the scheduling and routing problem separately. After implementing an initial version of each model, it would then allow us to easily add or alter constraints in the model. It would also allow us to have some freedom to experiment with the way in which the two separate models would interact. The downside of this method however is that the model is solved using an exact solver, in this case Gurobi [123], which is not fine-tuned to the characteristics of the (E)VRP. As a consequence, this only allows us to solve small problem instances. In theory this does not need to be a limitation, but while studying the performance of the model, larger problem instances and more complex objective functions allow for more interesting conclusions.

The second option we considered was directly working with OHD, ORTEC for Home Delivery. This is a really powerful engine that ORTEC uses to solve real customer home delivery problems. Using this would allow us to directly work with customer data. Because of that, it would give the most realistic outcomes, meaning that we could directly see the true potential of the new algorithm. The price of a model that can solve different problem variations all very efficiently is that it is very layered and complicated. It is therefore not straightforward how to add new functionalities or subproblems and apply them correctly. As it is crucial to be able to experiment with the model with at least some ease, we find that this option is not ideal for our application.

Another possibility is using an open source state-of-the-art Python-based VRP model as a starting point for the routing portion of the algorithm. An example of such a program is PyVRP [124]. While it is unlikely to tackle problems with the same level of detail as OHD is able to, it is able to solve larger instances than what an exact solver is capable of. Using Python for the scheduling part might also allow for better performance and more freedom in implementing additional heuristics. Adding features or experimenting with different formulations however might not be as trivial as it is when using AIMMS, but it is also nowhere near as complicated as it would have been when using OHD.

While using PyVRP seems like the most balanced option, we decided to use AIMMS for the entirety of the project. This is largely done due to practical reasons: the original model was created within AIMMS. This does not only include the technical part of the model, but also contains an application that allows for convenient initialization of different instances but more importantly displays the results in a pleasant and organized manner. Moving towards a different platform would be a large time investment, which is not guaranteed to pay off. The main downside of the use of AIMMS remains the limited input size for which a reasonable feasible solution can be found. While the performance of PyVRP might not be attained, it is possible to improve the performance of the AIMMS model by pre-processing the model beforehand. Therefore, we expect that this will not keep us from drawing interesting conclusions.

6.2. Client Data and Interests

While we are not collaborating with any clients directly, we did receive datasets from one of ORTEC's clients that makes use of the OHD software. Due to privacy reasons we cannot provide their name, so from now on they will be referred to as 'the company'. These datasets contain information about their depots, customers and vehicles. Half of these datasets contain tasks that need to be completed in the morning, the tasks in the other datasets need to be completed in the afternoon. These shifts can be fulfilled by either EVs, or conventional vehicles, both of which appear in the dataset. The EVs are only allowed to visit customers in the city center, while the conventional vehicles exclusively visit customers outside of the city center. This policy ensures that customers in the city center can always be visited, even last-minute, as conventional vehicles are not allowed in the city center. Additionally, this causes the EVs to only drive small trips and not run out of battery until the second shift of the day is completed.

When EVs are only a small part of the fleet, this policy is sufficient. If the share of EVs in the fleet increases, assuming that the geographical distribution of the customers remains roughly identical, this policy does become more troublesome. A solution would be to expand the set of customers that the EVs can visit. As a consequence, the EVs will no longer be forced to exclusively drive very short shifts, and therefore some kind of battery management is necessary. This could mean abiding by the battery restrictions in some other way, for example by enforcing a maximal shift length. It is also possible to improve performance by charging vehicles mid-day, which allows longer distances to be driven by the EVs.

We know that the company is looking to increase the number of EVs in their fleet, and at some point needs to let go of their current EV policy. If the EVs are only allowed to drive up to half of their range during a single shift, which is what happens if a vehicle is needed for two shifts a day without charging in between, they are severely limited in the set customers that they are able to visit. Allowing vehicles to charge between tasks would solve this issue, but this is no simple task when many vehicles but only a limited number of chargers are available. For applications such as this, the design of a charging schedule can greatly improve the ways in which the EVs are employed.

6.3. Model Objective Difficulties

A standard objective function for optimization models for the industry is heavily cost-based. That is completely expected, but simply using cost as the main objective function is insufficient for this application. That is because generally EVs are often still more expensive, in which case such a model will choose an ICE vehicle over an electric one. This is not even the main disadvantage: even at a competitive price point, EVs have a limited range and require significant down-time to refuel, which can be challenging when the vehicles are expected to be on the road for most of the day.

The datasets that we have access to already contain a breakdown of the costs attached to the vehicles. These costs are unknown for the vehicles the company intends to purchase, but based on the different costs attached to the available EVs, and the purchasing cost of the new EVs, we are able to make estimations. Despite that, we know very little about the meaning of these individual costs, and must therefore be careful when drawing conclusions. The comparisons between fleets containing existing vehicles and fleets containing new vehicles should in particular not focus on the financial picture.

6.4. Our Model Objective

This variation of the EVRP arose from a desire to make more use of an existing fleet of EVs. The primary model objective should therefore be a measure related to the occupancy of the EVs. There can however be secondary objectives, that can relate to the cost. A possible secondary objective could be to charge above what is needed (but below what is allowed) as much as possible, providing the driver with as much leniency as possible. Assuming the batteries do not discharge while not in use, such an objective does not lead to increased energy use. To acknowledge the cost of used energy or the used EVs in general does however contradict that previous objective. Other ways of extending the objective function of the model could be adding battery degradation (or other heterogeneous costs attached to the use of different chargers) or adding a cost for every time a vehicle is connected to, or disconnected from a charger. Different ways of extending the model are to introduce the option of adding new vehicles or chargers at a cost. This would then allow the model to be used as a tool to determine which vehicles and chargers need to be acquired such that the demand can be met, while the costs are kept as low as possible.

Summarizing this gives the following lists of terms that we assume are part of the objective functions for both the routing and the charging model:

Routing Model (Minimization problem)

- The number of kilometers driven $\times$ cost per kilometer
- The vehicles that are used $\times$ cost per vehicle
- The trips taken by each vehicle $\times$ cost per vehicle
- The number of vehicle uplifts for each type $\times$ penalty per vehicle uplift per type
- The number of customers that do not get visited $\times$ penalty per customer
- The number of time-windows missed $\times$ penalty per time-window
- The amount of energy in kWh exceeding the battery levels of the used vehicles $\times$ the price per kWh of additional energy used

Charging Model (Maximization problem)

- The number of shifts successfully scheduled $\times$ the incentive per shift
- The amount of energy in kWh contained in each of the batteries $\times$ incentive per kWh
- The amount of energy used in kWh $\times$ price per kWh
- The number of charging (dis)connections $\times$ penalty per charger (dis)connection
- The number of charger uplifts for each type $\times$ penalty per charger uplift per type
- The number of vehicle uplifts for each type $\times$ penalty per vehicle uplift per type

Note that not all costs, penalties and incentives need to be non-zero. As a matter of fact, we can assume that either the second or the third term of the charging model should have no effect on the objective, as they essentially contradict each other.

7 Modeling Details

7.1. Charging Model

This section describes the first part of the complete charging-routing algorithm. We start by providing the relevant notation and then give the initial model. This model assumes that each vehicle drove a shift, and returns to the depot at a fixed time with a known battery level. This depot contains a set of chargers that can each have a different charging speed. There is also another set of shifts planned for later that day, each of those shifts requiring a certain amount of energy. This model then aims to assign as many shifts as possible to the available vehicles. In order to be assigned to a shift, a vehicle needs to have been sufficiently charged by the time this shift starts. The initial objective is to simply assign as many shifts as possible. Later on we will assume all shifts need to be scheduled, instead judging the quality of a charging schedule by the amount charged and the extra vehicles needed to achieve this.

As we saw in chapter 6, there are quite a few choices for terms in the objective function. Reducing the number of charger (dis)connections, range anxiety and electricity costs all require some modifications to the model. There are other extensions that do not relate directly to the objective function, such as non-linear charging functions. These extensions will be discussed after the initial model has been given.

7.1.1. Initial Model

The sets, parameters and decision variables necessary for the initial model formulation are the following:
Sets:

- Ch: Chargers
- T: Timesteps
- S: Shifts
- K: Vehicles

Parameters:

- YRE^k: Battery level of vehicle k after returning from the previous shift
- YMX^k: Maximal battery capacity of vehicle k
- CCT_c: Charger capacity per timestep in kWh
- $AST_{s,t}$: $\begin{cases} 1 & \text{If shift } s \in S \text{ is active at timestep } t \in T; \\ 0 & \text{otherwise.} \end{cases}$
- PST_t^k: $\begin{cases} 1 & \text{If vehicle } k \in K \text{ is driving a previous shift at timestep } t \in T; \\ 0 & \text{otherwise.} \end{cases}$
- YMN_s^k: Minimal battery level needed to drive shift s for vehicle k
- FTA_s: The timestep $t \in T$ that shift $s \in S$ starts

Decision Variables:

$$x_s^k := \begin{cases} 1 & \text{if vehicle } k \text{ is assigned to shift } s \in S; \\ 0 & \text{otherwise.} \end{cases}$$

$y_t^k \in \mathbb{R}_{\geq 0}:$ battery level of vehicle $k \in K$ at timestep $t \in T$.

$$a_{c,t}^k := \begin{cases} 1 & \text{if vehicle } k \text{ is assigned to charger } c \in Ch \text{ at timestep } t \in T; \\ 0 & \text{otherwise.} \end{cases}$$

$z_{c,t}^k \in \mathbb{R}_{\geq 0}:$ amount charged at charger $c \in Ch$ by vehicle $k \in K$ at timestep $t \in T$.

We can then provide the formulation:

$$\max \quad \sum_{s \in S} \sum_{k \in K} x_s^k \tag{7.1a}$$

s.t.

$$y_t^k = y_{t-1}^k + \sum_{c \in Ch} z_{c,t-1}^k \forall k \in K \quad \forall t \in T \backslash \{1\}, \tag{7.1b}$$

$$y_1^k = YRE^k \qquad \forall k \in K, \tag{7.1c}$$

$$y_t^k \leq YMX^k \qquad \forall k \in K, \tag{7.1d}$$

$$z_{c,t}^k = a_{c,t}^k CCT_c \qquad \forall c \in Ch \quad \forall k \in K \quad \forall t \in T, \tag{7.1e}$$

$$\sum_{k \in K} a_{c,t}^k \leq 1 \qquad \forall c \in Ch \quad \forall t \in T, \tag{7.1f}$$

$$\sum_{c \in Ch} a_{c,t}^k + x_s^k AST_{s,t} + ST_t^k \leq 1 \qquad \forall k \in K \quad \forall t \in T, \tag{7.1g}$$

$$\sum_{c \in Ch} a_{c,t}^k \leq 1 \qquad \forall k \in K \quad \forall t \in T, \tag{7.1h}$$

$$x_s^k YMN_s^k \leq y_{FTA_s}^k \qquad \forall k \in K \quad \forall s \in S, \tag{7.1i}$$

$$\sum_{k \in K} x_s^k \leq 1 \qquad \forall s \in S, \tag{7.1j}$$

$$\sum_{s \in S} x_s^k \leq 1 \qquad \forall k \in K, \tag{7.1k}$$

$$x_s^k \in \{0,1\} \qquad \forall k \in K \quad \forall s \in S, \tag{7.1l}$$

$$a_{c,t}^k \in \{0,1\} \qquad \forall c \in Ch \quad \forall k \in K \quad \forall t \in T, \tag{7.1m}$$

$$y_t^k \geq 0 \qquad \forall k \in K \quad \forall t \in T, \tag{7.1n}$$

$$z_{c,t}^k \geq 0 \qquad \forall c \in Ch \quad \forall k \in K \quad \forall t \in T \tag{7.1o}$$

The initial objective function (7.1) simply aims to assign as many vehicles to shifts as possible. Constraints (7.1b) make sure that the battery level of a vehicle at a certain timestep (excluding the first) is equal to the battery level of the previous timestep together with the amount that battery was charged in that timestep. Constraints (7.1c) enforce that the battery level of a vehicle starts at the level it has upon return from its morning shift. The next constraints (7.1d) limit the amount of charge that a battery can hold. This value is determined both by a SOH policy determining a maximal SoC, and the battery capacity. Charging the battery is regulated by constraints (7.1e), determining that a battery can only be charged whenever it is connected to a charger, at a speed limited by the capacity of the charger. The next constraints (7.1f) restrict that each charger can only be connected to a single vehicle at every timestep. Constraints (7.1g) make sure that a vehicle can only be assigned to a single task. Either, it is connected to a charger, or it is driving either the morning or the evening shift. Next constraints (7.1h) avoid the situation in which one vehicle is charged by multiple chargers simultaneously. To assign shifts to vehicles, the vehicles need to have enough battery left to complete the shift. This is taken care of by constraints (7.1i). As this model does not penalize unassigned shifts, and instead rewards shifts that do get assigned, it is crucial to make sure that multiple shifts do not get assigned to the same vehicle, or that multiple vehicles perform the same shift. Constraints (7.1l) and (7.1j) impose this. The final constraints (7.1l) - (7.1o) enforce the binary and non-negativity limitations placed upon the decision variables.

7.1.2. Model Extensions

This section will expand upon the objective function alterations and model extensions that have been designed. We start by discussing the two main choices of objectives: either minimizing the energy cost, or maximizing the charging output. To end up with a usable charging schedule, it is critical to limit the number charger (dis)connections, which is discussed after. Then, we will provide the adaptations required to implement non-linear charging. Finally, some other minor modifications are suggested.

Objective Function Choices

There can be a variety of criteria that determine a good charging schedule, beyond scheduling all the available shifts. The main choice to be made is, should vehicles be charged as much as possible, above what is needed to complete their assigned shifts, or as little as possible, so that the energy costs are minimized?

When costs are a main concern, what is optimal depends on the type of energy contract the fleet owner has. If using electricity at night is priced at an off-peak rate, the latter choice seems obvious. In the end, roughly the same amount of energy will be charged in both choices, the difference being the additional energy depleted when the vehicles are not in use, that differs depending on the SoC of the vehicle. If (possibly short-term) cost savings are not the primary concern, or energy is cheaper during for example peak sun-hours, it might be better to charge the vehicles more during the day. This has several benefits, as the vehicles will not return to the depot completely drained after their final shifts. This is beneficial for the battery, as it is known that fully emptying batteries is bad for their SOH, but also for the drivers and the logistical puzzle that is charging all vehicles overnight. Drivers are less likely to need to drive vehicles that are expected to be fully empty on return, reducing range anxiety that drivers might have when it is uncertain whether their vehicles will be able to finish the trip.

If the goal is to minimize energy costs between charging moments, this can be implemented by subtracting a term $\sum_{k\in K}\sum_{c\in Ch}\sum_{t\in T} z^k_{c,t}$ from the objective function, multiplied by the energy cost. Otherwise, the maximal use of the chargers can be incentivized by adding the term $\sum_{k\in K}\sum_{s\in S} y^k_{FTA_s}$ to the objective function, possibly multiplied by a weighing factor.

Reducing Charger (Dis)Connections

With typical choices for the objective function, the result is that the solution could require a vehicle to use a different charger, or disconnect entirely, at every available timestep. While this might result in a very successful charging schedule, it can hardly be implemented as is. One way of dealing with this is increasing the length of the used timesteps. This makes sure that chargers can be (dis)connected less often, thus resulting in a less messy schedule. The downside of this however is that quite a bit of possible detail, and possible good solutions, do get lost. There is also still no guarantee of coherence. Alternatively, we can quantify the number of charger (dis)connections and minimize this number directly. In practice, we will find that the quality of the schedules found with such an adaptation is quite good, only scheduling two charging sessions in a single interval if absolutely necessary. The way in which this is done is adding a penalty for every time that a vehicle is connected to, or disconnected from a charger. This could be implemented by adding the following variables:

$$u^k_{c,t} := \begin{cases} 1 & \text{if on the interval } [t, t+1] \text{ a charger } c \in Ch \text{ is disconnected from or connected to a vehicle } k \in K \text{ ;} \\ 0 & \text{otherwise.} \end{cases}$$

To make sure that these variables behave as expected, we will define T' to be the set of timesteps excluding the final element and add the following constraints to the model:

$$u^k_{c,t} \geq a^k_{c,t} - a^k_{c,t+1} \qquad \forall c \in Ch \quad \forall t \in T' \quad \forall k \in K, \tag{7.2}$$

$$u^k_{c,t} \geq -a^k_{c,t} + a^k_{c,t+1} \qquad \forall c \in Ch \quad \forall t \in T' \quad \forall k \in K. \tag{7.3}$$

Constraints (7.2) and (7.3) make sure that when $a^k_{c,t}$ and $a^k_{c,t+1}$ take different values, $u^k_{c,t}$ cannot be 0. To be exact, constraints (7.2) take care of vehicles disconnecting from chargers, while constraints (7.3) attend to vehicles connecting to chargers. It is also possible to only implement a single one of these constraints, penalizing either connecting or disconnecting.

In order to minimize the number of charger (dis)connections, the value of these variables needs to be minimized, which is done by adding their sum $\sum_{c\in Ch}\sum_{t\in T'}\sum_{k\in K} u^k_{c,t}$ to the objective function.

Non-linear Charging

In section 4.2 we discussed the charging behavior of EV batteries. While this process is very non-linear, we did see that it could be approximated fairly well with a piecewise linear function. Our initial model, like most EVRPs found in the literature, uses a linear term to model the charging process. To improve accuracy, we could extend the initial model by instead using a non-linear charging function. Based on the work of Montoya et al. [100], we will choose two breakpoints for each charging speed. This could be simplified to two breakpoints expressed as fractions of the battery capacity. Each of the three resulting sections has a different charging speed for each of the charger types. Using this, the following sets and parameters can be defined.

- B: The set of battery level intervals differentiating charging speed. Subsequently define $B_{1,2}$ and $B_{2,3}$ to respectively be the first and second interval, and the second and third interval.
- $CCI^b_{NL,c}$: The charging speed at each interval $b \in B$ for every charger $c \in Ch$
- YMI^k_b: The maximal charge of the battery per section $b \in B$ for each vehicle $k \in K$

$$\omega_t^{k,b} := \begin{cases} 1 & \text{if the battery level of vehicle } k \in K \text{ is in interval } b \in B \text{ at timestep } t \in T; \\ 0 & \text{otherwise.} \end{cases}$$

$$\alpha_t^{k,b} \in [0,1] : \text{rational variable to determine the correct charging interval}$$

$$\gamma_{t,c}^{k,b} := \begin{cases} 1 & \text{if the battery level of vehicle } k \in K \text{ is in interval } b \in B \text{ and the vehicle} \\ & \text{is connected to charger } c \in Ch \text{ at timestep } t \in T; \\ 0 & \text{otherwise.} \end{cases}$$

To make sure that these variables take the correct values, we need to add the following constraints:

$$y_t^k = \sum_{b\in B} \alpha_t^{k,b} YMI_b^k \qquad \forall k \in K \quad \forall t \in T \tag{7.4}$$

$$\omega_t^{k,b} \le \alpha_t^{k,b} \qquad \forall k \in K \quad \forall t \in T \quad \forall b \in B_{1,2} \tag{7.5}$$

$$\alpha_t^{k,b} \le \omega_t^{k,b-1} \qquad \forall k \in K \quad \forall t \in T \quad \forall b \in B_{2,3} \tag{7.6}$$

$$\gamma_{t,c}^{k,b} \le \omega_t^{k,b} \qquad \forall k \in K \quad \forall t \in T \quad \forall b \in B \tag{7.7}$$

$$\gamma_{t,c}^{k,b} \le a_{c,t}^k \qquad \forall k \in K \quad \forall t \in T \quad \forall b \in B \tag{7.8}$$

$$\gamma_{t,c}^{k,b} \ge \omega_t^{k,b} + a_{c,t}^k - 1 \qquad \forall k \in K \quad \forall t \in T \quad \forall b \in B \tag{7.9}$$

$$\sum_{b\in B} \omega_t^{k,b} = 1 \qquad \forall k \in K \quad \forall t \in T \tag{7.10}$$

$$ca^k_{c,t} = \sum_b \gamma^{k,b}_{t,c} CCI^b_{NL,c} \qquad \forall k \in K \quad \forall t \in T \tag{7.11}$$

Each of these are additional constraints, except for constraints (7.11) that replace constraints (7.1e). Constraints (7.4) determine the values for $\alpha^{k,b}_t$. As $\alpha^{k,1}_t = 1$ will cause a larger model improvement than $\alpha^{k,2}_t = 1$ and $\alpha^{k,3}_t = 1$ respectively (this respectively implies that $\omega^{k,1}_t = 1$, $\omega^{k,2}_t = 1$ and $\omega^{k,3}_t = 1$ which in turn implies that might be possible to respectively set $\gamma^{k,1}_{t,c} = 1$, $\gamma^{k,2}_{t,c} = 1$, $\gamma^{k,3}_{t,c} = 1$, which implies that the amount a vehicle is charged might become non-zero), $\alpha^{k,b}_t$ will be set to the correct fraction. Constraints (7.5) enforce that $\omega^{k,b}_t$ can only be 1 if $\alpha^{k,b}_t$ is 1. Constraints (7.6) on the other hand enforce that $\omega^{k,b-1}_t$ will be 1 if $\alpha^{k,b}_t$ is 1. The next two sets of constraints (7.7) and (7.8) make sure that $\gamma^{k,b}_{t,c}$ can only equal 1 if both $\omega^{k,b}_t$ and $a^k_{c,t}$ are 1 too. To enforce that $\gamma^{k,b}_{t,c}$ will take value 1 in that scenario, constraint (7.9) is added. Constraint (7.10) additionally makes sure that the energy level of a battery can be only in one of the three different sections at the same time. Finally, constraint (7.11) ensures that the correct amount of energy in charged whenever a vehicle is attached to a charger, therefore replacing constraint (7.1e).

Limited Power Grid Capacity

When a large number of vehicles need to be charged simultaneously from the same location, it could be possible that the power demand is too high for the current power grid. It might not be possible to use all available chargers at the same time. In that case, the model can place a limit on the total power output at a given time. Defining CCT_{grid} to be the total available capacity of the power grid on an individual timestep, this can be done for example with the following constraints:

$$\sum_{c \in Ch} \sum_{k \in K} a^k_{c,t} \leq CCT_{grid} \qquad \forall t \in T \tag{7.12}$$

Constraints (7.12) are not relevant in every application. If all chargers can be occupied simultaneously then these constraints are clearly unnecessary. In addition to that, if using a faster charger is in no way more expensive than using a slower charger, and all-but-one of the available chargers can be used at the same time, then the same goal can be achieved by simply eliminating the slowest charger. When higher costs are attached to the use of faster chargers, this constraint does become relevant, as the model might prefer to use a slower charger instead of a faster one. If a moderately broad variety of chargers is available, the model could assign two slower chargers instead of a faster one, in turn allowing for more options than would have been available had we simply removed the option to charge from certain chargers.

Variable Electricity Pricing

There are many different ways in which the price for electricity can vary. There can be a differing day- and night-tariff, offering cheaper electricity during fixed nighttime hours. Or, the price of electricity might be lowest when the total yield of solar and wind energy is highest. An alternative is that the price of electricity does not depend on the time of use, but in part does depend on the maximal amount of electricity used simultaneously. It is also possible to pay a different tariff over the first however many kWh, than over the rest. While this last option will not impact the optimal charging schedule, the former two options could result in very different optimal schedules.

There are also other situations in which charging can be more or less costly depending on other circumstances. When the depot is generating renewable energy, for example by means of solar panels, a limited amount of energy could be considered free during peak sun hours. One might also charge a battery depreciation fee when using chargers with faster speeds. As faster charging, in particular DC fast charging, is known to reduce overall battery-life, it is not unreasonable to charge a small fee when using faster charging methods. Another way in which pricing can vary is when the purchasing costs of for example solar panels or faster charging methods are included, when the model is used to help determine if these investments will be worth it.

As long as the pricing depends linearly on the charger type and the time, one could subtract the following term $\sum_{c \in Ch} \sum_{t \in T} \sum_{k \in K} ECP_{c,t} z^k_{c,t}$ from the objective function. Here, $ECP_{c,t}$ is the price per kWh for each charger $c \in Ch$ at every timestep $t \in T$. When the maximal amount of electricity used simultaneously becomes relevant, one needs a different approach. Instead of limiting the power grid capacity as done by constraints (7.12), one could introduce a new variable that would take the place of the total available capacity. This new variable will then need to be minimized in the objective function.

Battery Degradation
We already discussed one way of reducing the battery degradation of a vehicle. Apart from the adaptation of asking a higher cost for using a faster charger, and the policy of only charging up to a certain percentage, one could implement one more measure. Despite the longer charging time and increased battery degradation of the final 20% we saw in section 4.3, fleet owners could want to charge past this 80% mark because this allows a vehicle to visit more customers. In order to allow this, while being aware of the consequences of doing so, one could charge an additional cost for charging past a fixed battery percentage.

In the previous section, we defined two breakpoints that in turn define three different charging intervals. If either of these (or possibly both) breakpoints is chosen as the moment from which on a small penalty is paid, then this is simple to implement when the model assumes non-linear charging. We would do this by adding a term $\sum_{c\in Ch}\sum_{k\in K}\sum_{t\in T}\gamma_{t,c}^{k,b}$ for some $b \in B_{2,3}$ multiplied by a penalty to the objective function.

Vehicle-Route Restrictions
It could happen that some vehicles are not allowed to travel certain routes, due to restrictions in size or emissions. The latter will of course mostly apply to ICE vehicles, but one can think of other reasons such a restriction might be necessary. An example could be vehicles that have a different load capacity: shifts designed with a larger vehicle in mind might not be able to be completed by smaller vehicles. Implementing such a restriction can be done by changing the value of a single set of parameters. Parameters YMK_s^k, that determine the minimal amount of energy that is needed for a vehicle k to drive a shift s, can be altered such that vehicles are unable to be assigned to a shift. We do this by setting the the battery requirement above the maximal battery capacity for that vehicle.

7.2. Routing Model

In this section the second part of the charging-routing scheme is is described. Since routing models for EVs were discussed in detail in the literature review, we will not provide any further explanation for the given model formulation. It should be noted that this model formulation originally allowed vehicles to visit charging stations between visiting customers, something that is not allowed by our model assumptions. For that reason, we only consider sets consisting of customer and depot nodes, as well as leave out constraints regulating the visits to charging stations.

One thing to keep in mind is that in some instances, we also allow the model to use ICEVs instead of only EVs. The main reason for this is that the company currently makes use of a fleet that consists of both EVs and ICEVs, and mimicking this is necessary to obtain a fair comparison. Using only EVs to deliver to a set of customers that are spread out widely will yield much worse routing outcomes than when a limited amount of ICEVs are also used, to visit the furthest customers for example.

One can think of multiple ways of involving ICEVs in a model designed for EVs, especially as the focus lies on the use of EVs. A few options are the following:

1. Consider a fleet only consisting of EVs. Whenever an EV cannot be assigned to a shift, an ICE vehicle will drive that shift.
2. First schedule all the available EVs, and only afterwards schedule a set of ICE vehicles to visit all the remaining customers.
3. Consider a mixed fleet of EVs and ICE vehicles, penalizing every ICE vehicle (or another metric quantifying the use of an ICE vehicle) that the model requires in order to schedule every single customer.

This first option is the simplest in regards to the EVRP formulation, although will be tricky to implement in practice. Either we allow all vehicles to drive a small distance only, resulting in the poor solutions predicted earlier. Or, we might allow the vehicles to drive longer distances, but risk the final solution almost exclusively containing shifts that are too long to be driven by EVs. For that reason, this method is insufficient. The second option requires the implementation of a second VRP model, but does result in maximal EV use. The downside is however that the customers that are not visited by EVs might lie in inconvenient locations, resulting in more kilometers driven overall due to a sub-optimal global solution. The third method uses only a single model, but it is more complicated than the standard EVRP model. Even if we simplify the ICE vehicle to an EV with practically infinite driving range, we still need to implement a suitable method for penalizing the use of those vehicles. Furthermore, the type and weight of the penalties need to be determined. The datasets we are using contain among other values, costs per used vehicle and driven kilometer by vehicle type. These costs are significantly higher for the ICEVs in the dataset, and therefore already act like suitable penalties. For that reason, this method is the most appropriate.

This model already contains most of the necessary features, as this initial model is not the most basic version of the EVRP. Despite that, there is a variety of alterations and extensions that were added to our working version of the model, such as the penalties necessary when adding ICEVs to the fleet. To mimic other features of the ORTEC routing algorithm, we also added the possibilities for restricting the maximal shift duration, using multiple depots and allowing vehicles to take multiple trips. We also adapted the model so that even when an instance turned out to be infeasible, a mostly correct solution would still be returned.

7.2.1. Initial Model

For the initial model, we simply used the Capacitated EVRP with Time Windows provided in section 3.2.2, leaving out the constraints related to charging stations. Because of this, we do not need to provide new notation, although for convenience we will repeat the used decision variables before giving the complete formulation:

$$x_{ij}^k := \begin{cases} 1 & \text{if vehicle } k \text{ travels from node } i \text{ to node } j,\ LEN_{ij} > 0; \\ 0 & \text{otherwise.} \end{cases}$$

$$y_i^k \in \mathbb{Q}_{\geq 0}: \quad \text{battery level of vehicle } k \in K \text{ on arriving at node } i.$$

$$s_i \in \mathbb{Q}_{\geq 0}: \quad \text{start time of service at node } i \in \mathcal{C}'_{0,N+1}.$$

$$\min_x \sum_{i \in \mathcal{C}_0} \sum_{j \in \mathcal{C}_{N+1}} \sum_{k \in K} LEN_{ij} x_{ij}^k \tag{7.13a}$$

s.t.

$$\sum_{j \in \mathcal{C}_{N+1}} \sum_{k \in K} x_{ij}^k = 1 \qquad \forall i \in \mathcal{C}, \tag{7.13b}$$

$$\sum_{j \in \mathcal{C}} x_{0j}^k \leq 1 \qquad \forall k \in K, \tag{7.13c}$$

$$\sum_{i \in \mathcal{C}_0} x_{ij}^k = \sum_{i \in \mathcal{C}_{N+1}} x_{ji}^k \qquad \forall j \in \mathcal{C}, \forall k \in K, \tag{7.13d}$$

$$y_j^k \leq y_i^k - (ECR^k LEN_{ij}) x_{ij}^k + VBC^k (1 - x_{ij}^k) \qquad \forall i \in \mathcal{C}, \forall j \in \mathcal{C}_{N+1}, \forall k \in K, \tag{7.13e}$$

$$y_0^k \leq VBC^k \qquad \forall k \in K, \tag{7.13f}$$

$$\sum_{i \in \mathcal{C}} \sum_{j \in \mathcal{C}_{N+1}} DEM_i x_{ij}^k \leq CAP^k \qquad \forall k \in K, \tag{7.13g}$$

$$s_i + (TRT_{ij} + SET_i) \sum_{k \in K} x_{ij}^k \leq s_j + LST_0 (1 - \sum_{k \in K} x_{ij}^k) \qquad \forall i \in \mathcal{C}_0, \forall j \in \mathcal{C}_{N+1}, \tag{7.13h}$$

$$EST_i \leq s_i \leq LST_i \qquad \forall i \in \mathcal{C}_{0,N+1}, \tag{7.13i}$$

$$x_{ij}^k \in \{0, 1\} \qquad \forall i, j \in \mathcal{C}_{0,N+1} (i \neq j), \forall k \in K, \tag{7.13j}$$

$$y_j^k \geq 0 \qquad \forall j \in \mathcal{C}_0, \forall k \in K. \tag{7.13k}$$

It should be noted however that in order to move from a homogeneous fleet to a mixed fleet, three parameters have been vectorized. Each vehicle $k \in K$ can now have a unique energy consumption rate ECR^k, battery capacity VBC^k and load capacity CAP^k. Another thing to keep in mind is that in later applications, the vehicles may start their shifts with partially charged batteries. In that case, parameter VBC^k could be replaced with a parameter VBL^k to denote the value of the battery level of that vehicle. This distinction is not important currently, but it does become relevant in the next section.

7.2.2. Model Extensions

There are quite a few features that need to be implemented before this routing model can provide solutions that are even moderately comparable with real routing schedules. For starters, to mimic the cost structure of the real routing solutions, we need to introduce the ability to pay a cost for every kilometer driven, vehicle used and customer visited. These functionalities can then be used to penalize non-electrical vehicles. In real routing scenario's, generated for the company, shifts can only have a certain duration. Vehicles can also be stationed at different depots, might not be allowed to visit every single customer location, and might be able to drive multiple trips during the same shift. These features were all important to implement, and the way this has been done is described below. Finally, we explain how we ensured getting a feasible solution, even if not all of the constraints are fully satisfied. This will not be applied in real settings, but these outcomes can be useful as intermediary solutions that have potential to improve the quality of the actual outcomes.

Penalizing ICE Vehicle Usage

There are three methods that we will use to penalize the use of non-electric vehicles. The first of these is to penalize the scheduling of such a vehicle. So, if an ICE vehicle is scheduled, we pay the penalty. We can also penalize the number of kilometers driven by such a vehicle, which acts as a stronger incentive to use the available electric vehicles as much as possible. A final method is to penalize the number of customers visited by the ICE vehicles, which has a similar effect.

To assign a penalty for every ICE vehicle that is used in the model outcome, one does not need to introduce any new constraints. Instead, we find that introducing a new variable and adding a term to the objective function is sufficient. This new variable is allowed to take binary values, and determines whether or not a vehicle is used:

$$v^k := \begin{cases} 1 & \text{if vehicle } k \text{ is used;} \\ 0 & \text{otherwise.} \end{cases}$$

We also need to introduce a new parameter. Instead of only focusing on the difference between EVs and ICEVs, we define this parameter in a more general way. This is as other vehicles might also be assigned a fee to be scheduled, which is modeled in the exact same way as an ICEV penalty. This parameter is defined as follows:

- RCV^k: The routing cost of vehicle $k \in K$

To make sure that the variable properly affects the rest of the model, we need to alter constraint (7.13c), which limits the number of times that a vehicle leaves the depot to 1. If we alter this constraint to instead say $\sum_{j \in \mathcal{C}} x_{0j}^k \leq v^k$ for every $k \in K$, we ensure that the vehicle can only leave the depot if this binary variable has the right value. Adding a term $\sum_{k \in K} RCV^k v^k$ to the objective function forces the model to only allow necessary vehicles in the solution.

Penalizing each kilometer driven by a single vehicle is slightly simpler, as it does not require modifying the set of constraints and variables that are used currently. The only thing we need to introduce is another parameter that contains the cost of each kilometer per vehicle:

- RCK^k: The routing cost per kilometer for vehicle $k \in K$

Note that here there is once again no structural difference between EVs and ICEVs; to differentiate between the two, different parameter values need to be chosen. To implement this feature, we need to multiply the objective function from the initial model formulation with this parameter. This results in the following new objective: $\sum_{j \in \mathcal{C}_{N+1}} \sum_{k \in K} RCK^k LEN_{ij} x_{ij}^k$.

To penalize the customers visited by non-electric, or any other set of vehicles, can be done similarly. We once more introduce a new parameter:

- RCT^k: The routing cost per task for vehicle $k \in K$

This functionality can be implemented by adding the term $\sum_{i \in \mathcal{C}_0} \sum_{j \in \mathcal{C}} \sum_{k \in K} RCT^k x_{ij}^k$ to the objective function. We see that this is correct, as the term $\sum_{i \in \mathcal{C}_0} x_{ij}^k$ expresses that a customer j is visited by vehicle k.

Maximal Shift Duration
Implementing the maximal duration of a single shift, while a little tricky, can be done with very minor model modifications. One might wonder however if it is necessary, seeing how EVs have a limited range, and cannot drive for a very long time without charging regardless. To account for unfortunately scheduled time-windows or the use of vehicles without a range restriction, we will include constraints that can enforce this. The parameter $MSD^k \; \forall k \in K$ will define the maximal duration of a shift driven by a vehicle. We can then define the following constraints:

$$
\begin{aligned}
s_i + SET_i + TRT_{i,EDV^k} x_{i,EDV^k}^k - s_j + TRT_{j,SDV^k} x_{j,SDV^k}^k \\
\leq MSD^k + 24 \cdot 60 \cdot (2 - x_{i,EDV^k}^k + x_{SDV^k,j}^k) \qquad \forall i,j \in \mathcal{C}, \forall k \in K
\end{aligned}
\tag{7.14}
$$

Because we only keep track of the moment a vehicle starts its service at a client, the formulation of this constraint is a bit more convoluted than one might have expected. In order to describe the duration of a shift, we start by finding the moment when the final service i of a vehicle has ended and has driven to the depot ($s_i SET_i + TRT_{i,EDV^k} x_{i,EDV^k}^k$). Next, we subtract the moment when the vehicle started driving, which is the starting time of the first service j minus the traveling time to that customer $s_j - TRT_{j,SDV^k} x_{j,SDV^k}^k$. Subtracting the second term from the first gives the complete duration, which we then proceed to upper bound by the given maximal duration. However, in advance we do not yet know which customers will be the first and last visited by a vehicle. Therefore, we only want to apply this constraint if both x_{i,EDV^k}^k and $x_{EDV^k,j}^k$ are non-zero. This is done by adding a term $24 \cdot 60(2 - x_{i,EDV^k}^k + x_{SDV^k,j}^k)$. If either, or both, of these variables are zero, then the fixed maximal shift duration is extended by an entire day (24 hours $\times$ 60 minutes) and essentially invalidates this constraint for those indices.

In section 3.2.2 we mention that duration time limits are one of the more common time-related restrictions. That however does not mean that such a constraint can be implemented in the same way in different problems. Lin et al. [125] for example also limited the total duration of a shift, by means of a more compactly formulated set of constraints. Formulating this restriction in that same way however does not work for our problem, because the existence of time-windows imply that sometimes the vehicles might be idle during their routes. The formulation by Lin et al. did not have this possibility, which simplifies the way that the total duration of a shift can be expressed.

Multiple Depots
To introduce multiple depots into the model, we need to re-evaluate the notation we introduced in the beginning. Initially, we assumed there to be a single depot, which appeared twice in the set of locations, in order to function as both the starting and ending point. To mimic this with multiple depots, we need introduce a set of depots D, in which each depot appears twice. We assume that we know what depot each vehicle starts and ends at: let $SDV^k \in D$ denote the starting depot of a vehicle $k \in K$, and $EDV^k \in D$ denote the ending depot of that vehicle. The sets of customers and depots now cannot be be written as $\mathcal{C}_{0,N+1}$, $\mathcal{C}_0$ and $\mathcal{C}_{N+1}$ anymore. We still assume the subscript to indicate the depot(s) that are contained in the set, resulting in expressions $\mathcal{C}_D$, $\mathcal{C}_{SDV^k}$ and $\mathcal{C}_{EDV^k}$.

With these changes, we rewrite the EVRP as follows:

$$\min_x \sum_{k\in K}\sum_{i\in\mathcal{C}_{SDV^k}}\sum_{j\in\mathcal{C}_{EDV^k}} LEN_{ij}x_{ij}^k \tag{7.15a}$$

s.t.

$$\sum_{k\in K}\sum_{j\in\mathcal{C}_{EDV^k}} x_{ij}^k = 1 \qquad \forall i\in\mathcal{C}, \tag{7.15b}$$

$$\sum_{j\in\mathcal{C}} x_{SDV^k j}^k \leq 1 \qquad \forall k\in K, \tag{7.15c}$$

$$\sum_{i\in\mathcal{C}_{SDV^k}} x_{ij}^k = \sum_{i\in\mathcal{C}_{EDV^k}} x_{ji}^k \qquad \forall j\in\mathcal{C},\forall k\in K, \tag{7.15d}$$

$$y_j^k \leq y_i^k - (ECR^k LEN_{ij})x_{ij}^k + VBC^k(1-x_{ij}^k) \qquad \forall i\in\mathcal{C},\forall k\in K,\forall j\in\mathcal{C}_{EDV^k}, \tag{7.15e}$$

$$y_{SDV^k}^k \leq VBC^k \qquad \forall k\in K, \tag{7.15f}$$

$$\sum_{i\in\mathcal{C}}\sum_{j\in\mathcal{C}_{EDV^k}} DEM_i x_{ij}^k \leq CAP^k \qquad \forall k\in K, \tag{7.15g}$$

$$s_i + (TRT_{ij}+SET_i)\sum_{k\in K} x_{ij}^k \leq s_j + LST_D(1-\sum_{k\in K} x_{ij}^k) \qquad \forall i\in\mathcal{C}_D,\forall j\in\mathcal{C}_D, \tag{7.15h}$$

$$EST_i \leq s_i \leq LST_i \qquad \forall i\in\mathcal{C}_D, \tag{7.15i}$$

$$x_{ij}^k \in \{0,1\} \qquad \forall i,j\in\mathcal{C}_D(i\neq j),\forall k\in K, \tag{7.15j}$$

$$y_j^k \geq 0 \qquad \forall k\in K,\forall j\in\mathcal{C}_{SDV^k}. \tag{7.15k}$$

Objective (7.15) only contains routes that start from the starting depot of a vehicle, and end at the ending depot of that vehicle. Constraints (7.15b) only force the vehicle to have arrived at their ending depot. The remaining constraints similarly replaced depot node 0 by the starting depot, and $N+1$ by the ending depot, for each vehicle. The term LST_0 in constraints (7.15h) was replaced by LST_D, as we assume that the opening and closing times of the depots are uniform. Without this assumption, these constraints need to be reformulated.

Limiting Vehicle Type per Customer
It might not be possible for each vehicle to visit every customer. This will mainly happen if conventional vehicles are no longer allowed in some city centers or other regions due to emission regulations. But it is also possible that goods that need to be delivered to a certain customer are not suitable to be delivered by certain vehicle types, for example when considering vehicle dimensions or (the absence of) temperature regulating features. To make sure that such limitations can be implemented, we need to introduce a binary parameter AVC_j^k for $k\in K$ and $j\in\mathcal{C}$ that takes value 1 if a vehicle is allowed to visit a customer, and 0 otherwise. We can then include the following constraint:

$$\sum_{i\in\mathcal{C}_{EDV^k}} x_{ij}^k \leq AVC_j^k \qquad \forall k\in K,\forall j\in\mathcal{C}. \tag{7.16}$$

Allowing Multiple Routes Per Vehicle
Commonly, vehicles may be expected to drive multiple routes in a single shift. This is particular likely to happen when the loading capacity of the vehicles is small, and can therefore only visit a small number of customers before needing to head back. For simplicity, our model did not contain this functionality initially. However, if we would like our model to be broadly applicable, this is a must. In the literature, this model variation is generally expressed using the Set Partitioning formulation (as described in section 3.2.1) of the VRP, such as done by Olivera and Viera [126]. As our model formulation is of a different type, we need to find a different approach. There seem to be two possible methods. The first method we considered is to allow the vehicles in the model to return to the depot multiple times before arriving at the ending depot.

This does require quite some restructuring efforts: Firstly we need to define a set of depots which the vehicles can visit to reload. These depots will need to be added as dummy depots, as many times as the maximal number of allowed trips. After visiting a reloading depot, the vehicle should be at its maximal loading capacity again. It is however not obvious how to implement this efficiently, if that is even possible. It appears that the model needs to keep track of which customer is visited during each trip in order to comply with the loading capacity constraints. As this is not an obvious task either, it is worth exploring another approach. The second possible method is to effectively add different copies of the vehicles to the model. This can be done by introducing an extra index to the variable x_{ij}^{k} so that it instead becomes x_{ij}^{kt}. Here, t is an element of the set of trips, which will be denoted as R. To not complicate the model too much, we will tightly limit the number of trips that vehicles can take, thus keeping the number of variables to a minimum. In particular, we define $MXT = |R|$. This adaptation, together with the introduction of those same dummy depots and one sets of constraints, are sufficient. These changes result in the following model formulation:

$$\min_{x} \quad \sum_{k\in K}\sum_{t\in R}\sum_{i\in \mathcal{C}_{SDV^k}}\sum_{j\in \mathcal{C}_{EDV^k}} LEN_{ij}x_{ij}^{kt} \tag{7.17a}$$

s.t.

$$\sum_{k\in K}\sum_{t\in R}\sum_{j\in \mathcal{C}_{EDV^k}} x_{ij}^{kt} = 1 \qquad \forall i \in \mathcal{C}, \tag{7.17b}$$

$$\sum_{j\in\mathcal{C}} x_{SDV^k j}^{k1} \leq 1 \qquad \forall k \in K, \tag{7.17c}$$

$$\sum_{i\in \mathcal{C}_{SDV^{kt}}} x_{ij}^{kt} = \sum_{i\in \mathcal{C}_{EDV^{kt}}} x_{ji}^{kt} \qquad \forall j \in \mathcal{C}, \forall k \in K, \forall t \in R, \tag{7.17d}$$

$$y_j^k \leq y_i^k - (ECR^k LEN_{ij})x_{ij}^{kt} + VBC^k(1 - x_{ij}^{kt}) \qquad \forall i \in \mathcal{C}_{SDV^kt}, \forall k \in K, \forall j \in \mathcal{C}_{EDV^kt}, \forall t \in R\text{'} \tag{7.17e}$$

$$y_{SDV^k}^k \leq VBC^k \qquad \forall k \in K, \tag{7.17f}$$

$$\sum_{i\in\mathcal{C}}\sum_{j\in \mathcal{C}_{EDV^{kt}}} DEM_i x_{ij}^{kt} \leq CAP^k \qquad \forall k \in K, \forall t \in R, \tag{7.17g}$$

$$s_i + (TRT_{ij} + SET_i)\sum_{k\in K} x_{ij}^{kt} \leq s_j + LST_D(1 - \sum_{k\in K} x_{ij}^{kt}) \qquad \forall i \in \mathcal{C}_D, \forall t \in R \forall j \in \mathcal{C}_D, \tag{7.17h}$$

$$EST_i \leq s_i \leq LST_i \qquad \forall i \in \mathcal{C}_D, \tag{7.17i}$$

$$\sum_{i\in\mathcal{C}} x_{ji}^{kt} \leq \sum_{i\in\mathcal{C}} x_{ij}^{k(t-1)} \qquad \forall j \in SVD^{kt}, \forall k \in K, \forall t \in R \setminus 1\text{'} \tag{7.17j}$$

$$x_{ij}^{kt} \in \{0,1\} \qquad \forall i,j \in \mathcal{C}_D (i \neq j), \forall k \in K, \forall t \in R\text{'} \tag{7.17k}$$

$$y_j^k \geq 0 \qquad \forall k \in K, \forall j \in \mathcal{C}_{SDV^k}. \tag{7.17l}$$

Here, the sets SDV^{kt} and EDV^{kt} respectively denote the starting and ending depots for each vehicle for each trip. For $t = 1$, $SDV^{kt} = SDV^k$ and EDV^{kt} contains EDV^k together with the first set of dummy depot nodes that vehicle k may visit. When $t = 2$ and above, (assuming $MXT > 2$), SDV^{kt} contains the $(t-1)$'th set of dummy depot nodes that vehicle k may visit, while EDV^{kt} still contains EDV^k together with the the t'th set of dummy depot nodes that vehicle k may visit. Finally, SDV^{kMXT} equals the $(MXT-1)$'th set of dummy depot nodes that vehicle k may visit and $EDV^{kMXT} = EDV^k$. When seen in the context of constraints (7.17d), we see that a trip may end either at the ending depot, or a dummy depot. If a vehicle arrives at an ending depot, it cannot make further trips (since it can only visit an ending depot once due to constraints (7.17b)), while arriving at a dummy depot allows another trip to start due to constraints (7.17j). As the vehicle capacity is limited per trip due to constraint (7.17g), visiting a dummy depot allows the vehicle to complete more deliveries, if the battery capacity allows it.

Allowing Incomplete routes

When routing vehicles that have a limited range, it can happen that the available vehicles cannot visit all customers in a single shift. The range may be too short, but it could also happen that the time windows are impossible to satisfy. Currently, if it turns out no solution exists that meets all the requirements, the model will become infeasible and won't provide any intermediary solution. However, it might be desirable to find such an intermediary solution anyways. This could simply be because delivering to all-but-one customer is better than not delivering to any. It is also relevant for future extensions whenever this model would be applied iteratively. If a solution cannot be found in the first iteration, it does not mean that it cannot be found in a later iteration. In order to find such solutions, it is critical that the routing model can provide solutions that do not meet all requirements.

There are a few different ways to allow such solutions. One method is to allow a route to not contain every single customer. For each customer that does not get visited, a hefty penalty is added. This can be implemented by introducing a new variable and a new parameter:

$$m_i := \begin{cases} 1 & \text{if customer } i \in \mathcal{C} \text{ does not get visited;} \\ 0 & \text{otherwise.} \end{cases}$$

- PMT_i: The penalty charged for not scheduling the task to visit customer $i \in \mathcal{C}$

Then, we need to alter one set of constraints. In the standard model, constraints (7.15b) make sure that every customer gets visited. To allow the model to not visit a customer, we can add the new variable to the left side of these constraints

$$\sum_{k \in K} \sum_{j \in \mathcal{C}_{EDV^k}} x_{ij}^k + m_i = 1 \qquad \forall i \in \mathcal{C}. \tag{7.18}$$

If $m_i = 1$, this constraint is satisfied, despite not visiting customer i. No alterations need to be made to the other constraints. Replacing constraints (7.15b) by (7.18), and adding the term $\sum_{i \in \mathcal{C}} PMT_i m_i$ to the objective function thus has the intended consequences.

It could also be possible to loosen the time-window restriction for certain customers. This will not solve infeasibilities regarding the range of the vehicles, but depending on the situation it could be desirable to have a delivery arrive late or early, compared to the next day or not at all. Depending on the strictness of the time-windows, two approaches can be chosen. Either, the time-window restriction can be removed entirely, allowing the delivery to be scheduled at any moment during the day. Or, we can add penalties depending on the size of the deviation. For simplicity, we will only provide an implementation for the first approach.

For the first case, we require the introduction of a variable and a penalty constraint that allows us to omit the restrictions of constraints (7.15i) entirely:

$$o_i := \begin{cases} 1 & \text{if the time-window restrictions of customer } i \in \mathcal{C} \text{ can be ignored;} \\ 0 & \text{otherwise.} \end{cases}$$

- PMW_i: The penalty charged for not scheduling the task to visit customer $i \in \mathcal{C}$ in the provided time window

Constraints (7.15i) must then be replaced by the following:

$$EST_i(1 - o_i) \leq s_i \leq LST_i(1 - o_i) \quad \forall i \in \mathcal{C}_D. \tag{7.19}$$

To further allow this model to provide infeasible solutions that are more informative about the reason that an instance does not have any feasible solutions, we implemented one more measure. When vehicles do not start their shifts fully charged, it might be possible that their range is just slightly too small to complete an efficient route. Giving the model freedom to allow a vehicle more energy use (up to the vehicle capacity) for a penalty might turn infeasible instances into feasible instances, combined with an updated charging schedule.

To do this, we need to introduce a new variable that determines the extra energy that the vehicle has used above its provided SOC, and a parameter that appropriately penalizes this when unavoidable:

$b^k \in \mathbb{Q}_{\geq 0}$: the energy added on top of the existing level for a certain vehicle of vehicle $k \in K$.

- PER^k: The penalty charged for exceeding the provided SoC of the partially charged vehicle $k \in K$

To implement this, we need to alter constraint (7.13f) to say $y_0^k \leq VBL^k + o^k$. Constraint (7.13g) remains unchanged, meaning that the battery level of a vehicle will never exceed past its maximal capacity.

7.2.3. Pre-processing the model

When running larger instances of MILPs, it can be a good idea to prune variables wherever possible in advance. This simplifies the problem that is going to be solved. In practice, this is done by shrinking the index domain of the used constraints and variables. We decided to remove the following set of variables from our model:

- The arc variables x for which the distance between two non-depot locations is larger than a small factor above the mean of the distances from each location
- The arc variables x for which the distance is longer than the range of the corresponding vehicle
- The arc variables x ending at a location that a vehicle can never visit because the demand exceeds the vehicle capacity
- The arc variables x belonging to vehicles and pairs of location that cannot be visited if the vehicle only visits those two locations before needing to head back to the depot

To further shrink the size of the convex hull, we experimented with adding different constraints that could speed up the model. In practice, these extra constraints only seemed to slow the model down.

7.2.4. Post-processing the model

In the previous section, we showed how to restrict the total shift duration time for each vehicle. This could sadly not be done in a very concise manner, illustrating the secondary role that time has in this model formulation. As a consequence, there is no obvious manner in which to ensure that the resulting schedule is as efficient timing-wise as we might like. For a schedule to become usable, it might be necessary to modify the timing of some scheduled tasks, in particular to reduce the total duration of the shift. Another critical part of the schedules that has thus far been omitted are driver breaks. Scheduling breaks directly in the VRP is generally considered quite challenging, which is why this was left out of the routing model. Coelho et al. [127] try to solve one such problem. They mention a few different options of adding break-rules directly into the VRP formulation, but conclude that this is too challenging to solve, and heuristics or other methods are used instead. We use a similar approach, and add the breaks in after the routes have been created.

Shift Duration Reduction

While analyzing the schedules that this routing algorithm generated, we found that shifts typically take longer than necessary and often have idle time between customer visits. For longer instances this was less significant, as the total duration of the shift was limited, but even then we saw that it was still desirable to modify the starting times of some of the individual tasks in the shift. For this purpose, we implemented a two-part algorithm. The order in which these parts are applied differs between the different types of shift that we aim to schedule.

In this context, we mean to differentiate between 'early' shifts and 'late' shifts. The first shift of the day is considered an early shift, and the last shift of the day is considered a late shift. If there are more than two shifts during a single day, then either type can be assigned to the remaining shifts. The reason for making this characterization is that early shifts can have their starting time moved forward with little benefit, whereas moving back their ending time could result in more opportunity to charge the vehicles for the next shift. For late shifts, the opposite reasoning holds: moving the starting time forward will allow more time to charge the vehicles, while moving the ending time back is unlikely to result in a similar benefit.

To fully make use of these differences between shifts, we treat both types differently. We aim to schedule early shifts as early as possible, and late shifts as late as possible. This results in the most time in between these two shifts. This difference is made by the order in which the parts of the algorithm are applied. The first part of the algorithm tries to move each task forward as much as possible, while the second part moves all tasks backwards. A compromise is also possible, in which only the tasks that directly affect the starting or ending time of the shift are moved. The pseudocode of the algorithms is given below.

Part I

1. Either select:
   (a) The final task of the shift, or
   (b) The rightmost shift that when moved forward will no longer impact the ending time of the shift.
2. Move this task forward as much as possible, so that either it finishes at the end of its time-window, or right before the next activity has been scheduled.
3. As long as another task is scheduled before the current task, choose this task and repeat step 2.

Part II

1. Either select:
   (a) The first task of the shift, or
   (b) The leftmost shift that when moved back will no longer impact the starting time of the shift.
2. Move this task back as much as possible, so that either it starts at the beginning of its time-window, or right after the next activity has been scheduled.
3. As long as another task is scheduled after the current task, choose this task and repeat step 2.

For early shifts, we first run part II version (a) and then part I version (b). This results in an ending time that is as early as possible, while the remainder of the activities are scheduled as closely together as possible. For late shifts, we first run part I version (a) and then part II version (b). This time, the starting time is as late as possible, while the rest of the activities are again scheduled as closely together as possible. For shifts that are not clearly late or early, one can choose if starting earlier or later is preferred, and choose the pair of algorithms accordingly.

Break Insertion

When drivers are expected to drive long shifts, it is crucial that they can take a break every so often. This is not only important for driver well-being and safety, it is also a legal requirement. In the EU, drivers must have a 45 minute break every 4.5 hours, which can be split into a break of at least 15 minutes and a break of at least 30 minutes. On top of that, exceptions excluded, a driver may drive up to 9 hours a day. More information can be found in the Driving and Resting Time Rules of the European Labour Authority [128]. As argued before, we decided to not implement such break-rules directly into the formulation, and therefore have to add these breaks in heuristically afterwards. As there is typically some idle time between customer visits, and the time windows are not too short, we should always be able to add sufficient break time to the schedules.

There are many ways of inserting the breaks into our schedules, as there is no single optimal moment for each break to be scheduled. We will now describe the way we have approached this problem. First, we find the number of breaks that need to be inserted into every shift. This depends on the total duration time of each shift, which is only known after the shift duration reduction procedure has been applied. So, we apply this procedure, calculate the total duration, and based on that find the number of breaks needed. As this schedule contains less gaps than its predecessor, we will revert back to the initial schedule, and then insert the breaks into the gaps at appropriate moments. We then apply the shift duration reduction procedure again, now including the scheduled breaks. For our instances, this scheme has always functioned correctly, but it might happen that there are insufficient gaps in the original schedule. To make sure that we scheduled all necessary breaks, at the end of the procedure we will check that there are no breaks left to schedule based on the previous calculations. If this is not the case, we should still insert the missing breaks manually. The complete scheme is summarized below:

1. Calculate the number of breaks that need to be inserted per shift.
2. Before applying the shift duration reduction procedure, insert the breaks.
3. Apply the shift duration reduction procedure.
4. Check if the breaks have been scheduled correctly, if not, correct.

7.3. Model Variations

This section will describe the relationship between the charging and the routing model, and how the two interact. We first describe the simplest way these two algorithms can cooperate in order to determine the optimal way of charging the vehicles, based on the scheduled routes. The second variation is very similar, except this time we allow the schedules to be made back-to-back, and therefore incorporate information about the energy levels after returning from their previous shifts in the charging model. Keep in mind that we are still assuming the vehicles are fully charged before calculating the routes. If there is sufficient opportunity to charge, that assumption is reasonable, but otherwise we risk our routing model providing unrealistic routes. To reduce this risk, we will estimate the available energy in the vehicles after charging, before the routes have been made. This is done in variation 3. We also introduce a final model variation that closely resembles variation 2 and 3, except that this time we allow vehicle and charger uplifts.

7.3.1. Variation 1: Charging Scheduling Only

This first model variation calculates the routes and the charging schedule separately for only a single charging instance. First, the vehicles are scheduled to visit a set of customers, both for two shifts during the day. This gives us information about when the vehicles are available to charge, and how much energy is needed. Based on this information, a charging schedule is made. This scheme is illustrated in figure 7.1.

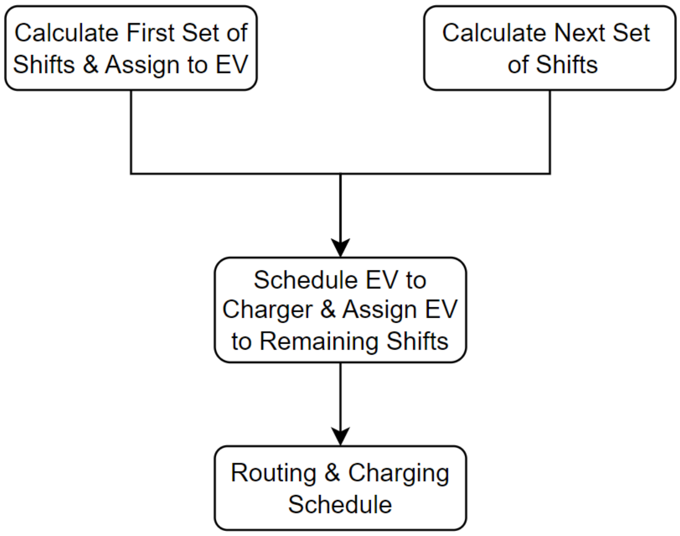

Figure 7.1: Flowchart of the first model variation.

This model does have limited use. If there is an abundance of chargers available, in particular when small instances are solved, finding a charging schedule is not a very challenging task. If on the other hand the number of chargers or the allowed charging time are restricted, it is likely that not all vehicles can be charged completely. In that case, it is not obvious which vehicles need to be charged, so that all shifts can be assigned to the vehicles. When the opportunity to charge vehicles shrinks even further, this problem becomes even more challenging, and requires the third variation to find feasible overall solutions.

7.3.2. Variation 2: Back-To-Back Routing and Charging Scheduling

This second model variation resembles the first in the sense that no functionalities have been added. The only difference is that this variation runs multiple instances back-to-back. This model variation can be applied when there are more than two shifts during a single day, or in order to make a suitable overnight charging schedule. The benefit of using this model variation instead of simply running the first variation multiple times, is that information about the SoC of the vehicle batteries and the vehicle-shift allocations are applied directly.

Due to the repetitive nature of this model variation, the flowchart of this model will naturally look slightly different to that of the previous variation. To illustrate the workings of this variation, the following flowchart is provided in figure 7.2 below.

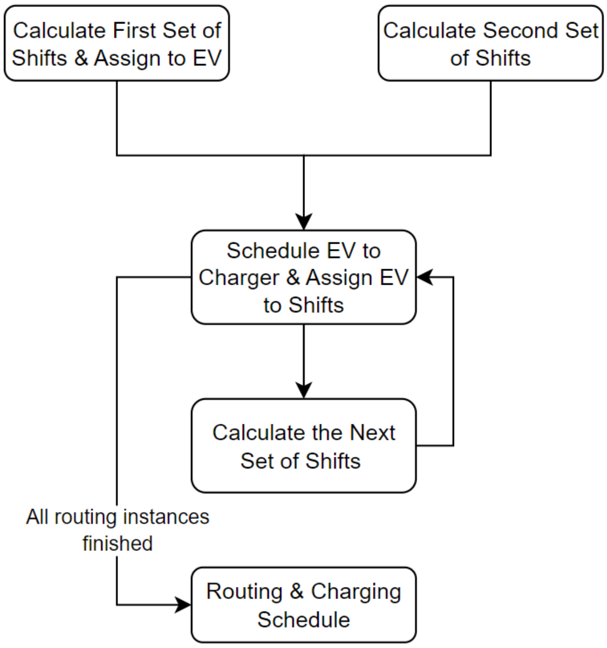

Figure 7.2: Flowchart of the second model variation.

7.3.3. Variation 3: Routing and Charging Scheduling with Preliminary Battery Estimates

As argued before, the solutions of the previous model variations will start to decrease in quality if there is insufficient opportunity to charge the vehicles. Knowledge about the expected SoC of the vehicle batteries can be incredibly valuable when calculating the routes these vehicles will have to drive. It can for example make the difference between a set of routes that are all slightly too long for the available vehicles, and hence to complete the shift an entirely new set of vehicles is needed, and a set of routes that uses an extra vehicle (or possibly more), but the majority of the vehicles that have driven a first shift can also drive a second shift.

The challenge with estimating the battery levels of the vehicles before a routing instance is that the battery available before driving a shift depends both on the outcome of the charging and the routing model. The starting time of the shift created by the routing model, assigned to a vehicle by the charging model, directly impacts how much the vehicle can be charged before the shift starts. We first provide the flowchart of this model variation in figure 7.3 to illustrate the altered scheme, and then explain the method used to approximate the SoC of the battery before starting a shift.

Preliminary Battery Estimation

Up until now, we have assumed that all the vehicles start their shift with a full battery. If it is the first shift of the day, we assumed that the vehicles have been charged completely overnight, but for later shifts that day, this assumption no longer holds up. One way of getting an accurate battery estimation is to run the routing algorithm under the assumption of fully charged vehicles, and then running the charging model. The battery levels that this model found can then be used to run the routing model again. This is however

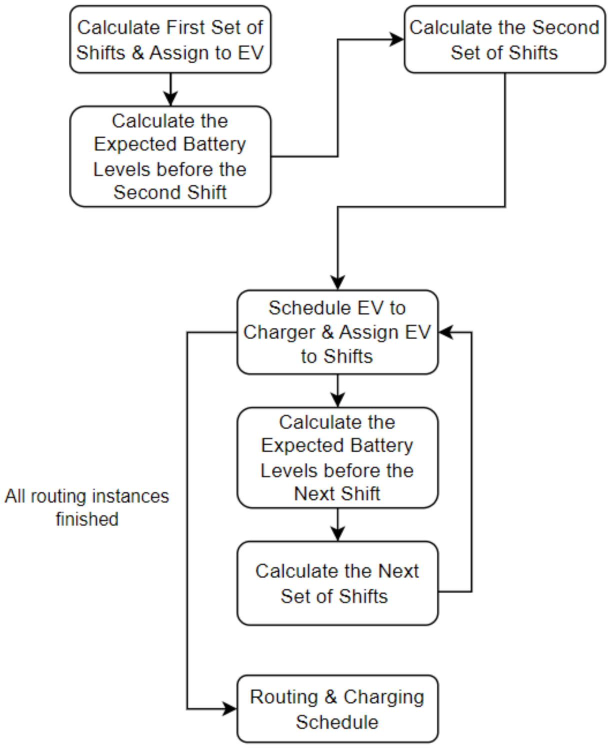

Figure 7.3: Flowchart of the third model variation.

very time consuming, as we run the routing and charging model twice for all but the first routing instance. There are a few alternative methods with which we can estimate the available energy before a shift in a more efficient manner:

1. Assume that we charge each vehicle using the fastest charger during a period of time that is at most as large as the time that vehicle will have available for charging.
2. Assume we charge each vehicle by the average total charger output. What we mean by this is that we assume that each vehicle is charged by each of the chargers the same amount as all other vehicles.
3. Run the version of the charging model that maximizes charging output before the routing model, without asking for a minimal amount of energy to charge a shift.

The first method is very simple, and mostly useful when there is little time available to charge, but no large shortage of chargers. The second method, the details of which can be determined in different ways, is most useful when the shifts are expected to have roughly equal length. The final method is more computationally expensive, but can result in the most sophisticated upper bound. For all of these methods it is crucial to be able to approximate the start times of the shifts as accurately as possible. When the available charging time is overestimated, we might as well have assumed that the vehicles were fully charged upon starting their second shifts. If we underestimate the time available to charge, we risk ending up with poor outcomes of the routing model. To avoid both of these scenarios, we will run a very simplified version of the routing model to gain some insight in the starting moments of the shifts. There is not a single way of doing this, and more experimentation might be necessary to find the most accurate results, but the routing instance we are running for this purpose operates under the following assumptions:

- Only consider the variables x_{ij}^k, s_i and possibly v^k.
- Only consider constraints (7.13b), (7.13c), (7.13d), (7.13h), (7.13i).
- The set of customers to visit is twice as large as the number of vehicles we expect to use, and consists of the customers that need to be visited first according to their time windows.

- Each vehicle in the set of vehicles expected to be used must visit at least one customer. We enforce this by adding a term to the objective function penalizing vehicles that did not leave the depot.

This means that we only calculate the routes for the customers that need to be visited earliest, omitting the range and demand restrictions from the model while making sure that all vehicles are used. This will result in estimates for when the actual routing model will schedule the start of the shifts. These estimates are then used as input for the preliminary charging model. The output of this model will finally yield the battery estimates for the vehicles before the real routing instance is solved.

7.3.4. Variation 4: Routing and Charging Scheduling with Vehicle and Charger Uplifts

A final variation of the presented models is the version that allows the vehicles and chargers to be uplifted from the currently available set of vehicles and chargers. This version is particularly useful when companies want to invest in new vehicles and/or chargers, but are uncertain about the details. The previous model versions can also be used for this purpose, when for example a set of vehicles and chargers is made available that is much larger than necessary. The set of used vehicles and chargers can then be seen as an approximation of what is needed. A downside of this method is that the cost of these investments is insufficiently considered. To solve this, a new model variation is proposed.

This variation can act as an extension of any of the previous three variations. To illustrate the flow of this variation, we chose to adapt the flowchart of the second model variation. This flowchart is given in figure 7.4 below.

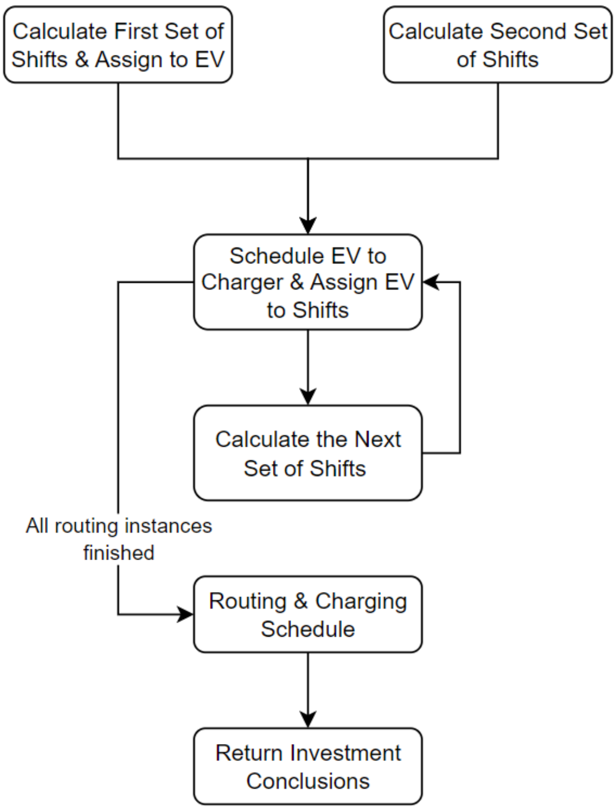

Figure 7.4: Flowchart of the fourth model variation.

To uplift vehicles, the same mechanism is used as presented in section 7.2.2. Instead of assigning a cost to ICEVs only, each vehicle k that can be uplifted is assigned a value RCV^k. This value can be (a fraction of) the purchasing price of the vehicle, but that is not necessary. As long as the uplifting costs for different vehicles are proportionate and not so high that the model prefers an infeasible solution over a feasible one containing uplifted vehicles, the exact numbers do not matter. Note that we can introduce this change in both the routing as well as the charging model. To uplift vehicles in the charging model, a new constraint is needed: $\sum_{s \in S} x_s^k \leq v^k$ for all vehicles k in the set of vehicles allowed to be uplifted. Just like in the routing model, we also add a term to the objective function that increases the objective by the uplift cost whenever v^k is non-zero. Note that v^k is both a variable that can be introduced into the routing and in the charging model.

For charger uplifts, the same idea is used. Here, we do introduce a new variable, that determines whether or not a charger is uplifted:

$$w_c := \begin{cases} 1 & \text{if charger } c \in Ch \text{ is used;} \\ 0 & \text{otherwise.} \end{cases}$$

To be able to use this variable, we need to adapt constraint (7.1h) to only allow a charger to charge a vehicle if it is either allowed or uplifted. This constraint will now look as follows: $\sum_{c \in Ch} a_{c,t}^k \leq w_c$ for all $k \in K$ and $t \in T$. We once more add a term to the objective function that takes the value of the charger uplifting cost whenever w_c is non-zero for a vehicle that was not already owned.

Part III

Results

8 Tests

8.1. Baseline Assumptions

In order to accurately test the performance of our model, we need to know how the instances are solved in practice. The data given by the customer contains only the load capacity of the vehicles, the battery capacity of the electric vehicles is not included. With the current assumptions however, it is known that the EVs will not run out of battery before arriving back at the depot. To be specific, the EVs used by this customer will only deliver to their clients in the city center, while the ICEVs will deliver to the clients outside of the city center. Due to the limited set of tasks in the city center, the EVs will automatically make only short trips, allowing the vehicles to drive two shifts without charging during the day. To fulfill the objective of our model, i.e. increasing the number of kilometers driven by an EV, we need to let go of this assumption. To fairly compare the two situations (before and after allowing vehicles to charge during the day), we will assume that EVs may visit customers outside of the city center. Now, it does become crucial to know more about the battery capacity and energy consumption of the vehicles used.

To the best of our knowledge, the EVs that are currently in use are a customized version of the type Volkswagen Crafter. The exact details of this customization are unknown to us, but what we do know is that the maximal range of this vehicle is $200km$. To be certain that the vehicles can complete their routes, in particular during tougher circumstances, we will divide this range by 1.5, resulting in a vehicle range of $133km$. This is a typical reduction of the range for daily use, suggested by for example the The Royal Dutch Touring Club ANWB [129].

It is also known that the customer is planning on expanding their fleet by purchasing new EVs of the type Toyota Proace Electric. There exist different editions of this vehicle, both in form factor as well as the size of the battery. The vehicle can either contain a $50kWh$, or a $75kWh$ battery, and as of yet it is unclear which of these variants, of what combination will be purchased. Additional uncertainty is added due to the custom form factor. The standard configurations have a WLTP [130] range of up to $230km$ for the smaller battery, and $330km$ for the larger one [131]. Due to the less aerodynamic shape of the customized vehicle and the added energy-use from a necessary freezer compartment, these estimates of the range are not fully accurate for the custom version of the vehicle. The WLTP estimates are based on ambient temperatures near the European average, which is similar to the climate in the Netherlands [132]. This, in addition to the fact that the Netherlands is very flat country, implies that the expected range of the vehicles will not decrease much further due to environmental factors. Many other factors that influence energy consumption, such as described in section 4.1.3, are highly variable or unknown to us, and therefore cannot be used to improve the range estimation. The driving profile used in the WLTP estimates is also quite varied, which does not allow us to make obvious corrections to the energy use. The only other factor that we can correct for is the vehicle load. While the vehicle load keeps decreasing while the vehicle completes more stops on its routes, for a large part of the trip the vehicle does contain a lot of extra weight. To fairly account for this extra weight, and to be sure that in less-than-ideal circumstances the vehicles can safely complete their route, we will similarly reduce this range, resulting in vehicle ranges of respectively $147km$ and $230km$. Something else we know about these new vehicles is that their length is almost one third smaller than the old ones. As this includes the length of the cabin, the loading capacity does not decrease proportionally to this decrease in length. This is the only information we have about the size difference between the old and new vehicles, so we will simply assume that the new vehicles have 50% of the loading capacity of the old vehicles.

8.2. Testing Methodology

In the previous section, the baseline assumptions were laid out. Using those assumptions, we can create different instances that reflect realistic settings against which we can compare the outcomes of our own model. These instances can then be solved using the OHD engine, which is how real instances of this sort, in particular the datasets that we are using here, can be solved. There are two key differences between the cases we can recreate, and the real cases. For starters, our exact routing model can only solve very small instances exactly, and only slightly larger instances to within an acceptable optimality gap in a realistic time-frame. Secondly, while our routing model does consider almost all of the information the dataset provided, it does not take congestion into account. Doing so would require the travel times to become time-dependent. This is impossible to achieve accurately with only minor modifications to the model formulations, so we decided to ignore congestion factors all-together. Our model objective also does not take the hourly rate of the drivers into account, for similar reasons. Luckily, this rate is uniform per vehicle type. Removing the hourly cost, replacing it with a fixed cost per vehicle, should result in an acceptable approximation for a schedule and cost-picture.

Due to both of these factors, we are inevitably dealing with simplified versions of real instances. While leaving out congestion is relatively inconsequential in the larger picture, the modification of only using a small subset of customer locations to visit is accompanied by several decisions. First of all, the size of the subsets need to be determined. Secondly, we need to determine which of the customers will be included in these subsets. Another question is what happens to the set of vehicles that can be used? Finally, we need to determine which chargers are available for the vehicles to use. These questions are answered in the next subsection.

8.2.1. Modifications related to instance-size

The first thing to determine when creating an instance is how many customers will need to be visited. When only a small number of customers are visited, the instances can be solved to optimality in a small amount of time, but are not very interesting to study and do not illustrate the full potential of the charging model. If the routing model cannot find a reasonable solution within a reasonable amount of time due to the size of the subset of customers to visit, then we cannot draw strong conclusions from those outputs. To balance these two points, we decided to pick an instance size of 50. We have not solved instances of this size to optimality, but after running these instances for up to two hours we did find solutions that were good enough. The gaps for these instances typically ranged from around 3% to 12%, which did not improve much after the first hour.

The next question to ask is which 50 customers to visit from each of the datasets. The difficulty of an instance is not only determined by the number of customers to visit, but also their placement on the map. The distance to the depot(s) and the amount of clustering can greatly impact how easily a solution can be found. The demand of each customer can also determine the difficulty of finding an optimal solution, although the default capacity of the vehicles is so large that this did never appeared as a bottleneck.

Because the OHD engine will solve the same instance as the exact routing algorithm, the difficulty of an instance does not impact the value of the comparison. Ideally however we wish for instances that somewhat reflect realistic settings, while not overburdening the routing model. For that reason, together with the limited amount of time and resources available, we will only test a single instance per instance size per dataset. Due to only minor importance of this selection, we will simply pick one of several uniformly sampled sets that contains at least one but less than 10% 'outlier' customer locations. By an outlier customer location we mean customers that are further than 40% of the range of the currently owned EV away from the main depot. In this case, that gives a distance of $54km$.

In regards to the available vehicles for specific instance-sizes, we can simply allow access to the entire available fleet. In the exact model, choosing an appropriate number of vehicles can benefit the performance of the model, but when the solutions are found using OHD, this is of no concern. Practically, this means that we have access to an abundance of EVs. As a consequence, due to their higher cost, ICE vehicles will only be used if the vehicle range is too small to efficiently reach certain customers. As long as we account for this behavior in the measure we use to determine performance, this is is preferable over artificially introducing conventional vehicles. While this might be more reflective of the current situation, we actually do not know what the appropriate proportion of EVs against ICEVs should be for our instances. Additionally, deliberately adding conventional vehicles to the situation simplifies the accompanying charging problem, which is undesirable in particular for the smaller instances.

8.2.2. Modifications related to vehicle policies

As mentioned previously, we need to adapt the assumptions under which the model is operating currently. This leaves some freedom to determine the new vehicle policies. The policy that most closely mimics the real policy is to allow the EVs to drive up to half of their vehicle range for each session, so that the vehicles can be equipped during both shifts. Alternatively, we can assume the vehicle drives as much as needed, and is left with an almost empty battery for the shifts that come after. It is possible to combine these policies, and split the allowed battery use between shifts in any other way imaginable, but this is unlikely to be successful, as the number of possible routes a vehicle can drive does not increase linearly with the allowed range.

Additionally, we know that this customer will have access to a larger and more varied fleet of EVs in the future. It would be a waste not to use this information in our calculations, outside of illustrating the vehicle investment functionality of the model. We know which vehicles will be purchased, but other than that we have no information. It is therefore not at all obvious how to combine the new vehicles with the current mixed fleet. A better solution is to keep the two separate, and run the same instances assuming we only have access to the new vehicles. Only when testing the vehicle investment functionalities of the model should we combine these two fleets.

There is also another question to be asked: what do we assume will be the starting battery level of our vehicles? It is simple to assume they are fully charged, in particular if the shifts start in the morning and the vehicles have had a chance to charge overnight. This is also necessary when we use model variation 1. When we want to run the instances back-to-back, we can use previous battery level information to complicate the charging problem. This can also give interesting results, as it might not be possible to charge all vehicles fully overnight, leaving the vehicles partially charged in the morning.

8.2.3. Remaining Parameters

In this section we broadly sketched the settings we would like to test. Some of the parameters that are needed for these instances, such as vehicle properties, have been determined in section 8.1. Others, such as the size of the instances, were chosen earlier in this section. There are also many parameters that we leave unaltered from the datasets, related to the cost of using the vehicles. We did remove the hourly rate from the dataset. To compensate for this, we did increase the cost per individual vehicle proportionally to the hourly rate assigned to that vehicle. A final parameter that has until now been left undiscussed is the price of energy per kWh. This price varies daily, even from moment to moment, so there is no single correct choice. To not complicate matters further, we will choose a fixed price of €$0,42$ per kWh, taken from [133] in November 2023.

Another set of parameters that is as of yet undetermined is the cost per driven vehicle for the fleet consisting of the two new types of EVs. In particular, we need to determine the cost of using each of those vehicles. For the outcome of the routing model, the exact values are not relevant, as long as they are sufficiently high and the difference in cost due to the different battery sizes between the two models is appropriate. For that reason, we can deduce a daily cost from of the listed price for the standard form factor for both vehicles. The listing price for the $50kWh$ model of the Toyota Proace Electric is €40.450, while the $75kWh$ version costs €47.900. Our cost approximations for daily use might not reflect their true expense, but will likely be proportioned reasonable correctly.

Mostly ignored thus far is the available charging infrastructure. We have no information about the available chargers at the company, nor if they intend to extend their set of chargers, so we are left to make assumptions. To us, it seems like a good estimate to have access to a total number of slower chargers that is about equal to half of the vehicles we are expected to be used, and to add a set of fast-chargers that can charge at least 10% of the vehicles simultaneously. This allows the company sufficient charging opportunity to charge the vehicles overnight, and to charge some vehicles quickly if necessary. Another point to consider is that chargers are typically installed in pairs, so for each type of charger installed, an even number is preferred. The number of chargers needed will differ on a case-by-case basis, as we are testing a variety of instances. As for all instances we have provided a baseline, the number of EVs used in these baseline solutions will be used to determine the set of chargers we have access to.

8.2.4. Scenarios

Ideally, we would consider every possible scenario for each dataset and all instance-sizes. Running a routing instance, in particular of a larger size, can take a long time. Thus, we need to be selective about which instances to use in our tests. To create an instance, we need to answer each of the following three questions:

1. Which instance size is used? 50 customers, or more, if our routing algorithm does not need to solve this instance?
2. Which vehicle fleet will be used? Do we choose a mixed fleet of conventional and currently owned electrical vehicles, a fully electric fleet consisting of two new types of EVs, or a mixture of the two?
3. Which vehicle policy do we assume is used?

(1) Each routing instance, we allow half of the battery to be used. Then, it is guaranteed that each vehicle can drive two shifts each day.

(2) Each routing instance, the vehicles are allowed to use their entire battery. If they cannot drive another shift the next routing instance without charging beforehand, another vehicle will need to be scheduled.

(3) Each routing instance, the vehicles are allowed to use their entire battery. To make sure that these vehicles can complete other routes that same day, they will be charged before each starting their next routes.

In the next section the different cases we will test are laid out. For each of those cases, a suitable set of instances is created based on these options. In some cases, these instances are then solved under different, case-specific, assumptions.

8.3. Experimental Cases

As this model has different possible applications, in order to accurately test the performance we need to set up different experimental cases. For each case, different initial choices can be made. These cases and initial values are created somewhat artificially, as described in section 8.2.

Quickly summarized, this model has the following features:

1. Solving the Vehicle Routing Problem
2. Determining a charging schedule
3. Determining an optimal set of vehicles to acquire
4. Determining an optimal set of chargers to install

Features 1 and 2 can be considered as standalone models. Despite the goal of this project not being the design of a high-performance VRP model, the quality of the solutions should be evaluated. While for a routing model we require data that is typically not too hard to obtain, the charging model requires solutions from a routing algorithm. As our routing model currently forms the bottleneck of the model, ideally we would equip a different routing model in order to test the capabilities of the charging model on its own. While these two models can be executed separately, they can be combined, resulting in improved solutions for both models. Feature 3 can be applied when using either, or both, of the algorithms. To use feature 4 however, the charging model needs to be involved. Overall, we can examine the performance of these functionalities by studying four different cases:

1. Testing the VRP only
2. Testing the charging model with routing outcomes from OHD
3. Testing the combined vehicle routing and charging model
4. Testing the vehicle and charger investment functionalities

For each of these cases, we will discuss the goal of running the experiments, and which instances we would like to test.

8.3.1. Case I: VRP

The first set of tests we do is to evaluate the quality of the routing algorithm. As our algorithm is exact, the quality of the solutions could outperform the quality of the routes generated by OHD, the heuristic that is currently used to solve these instances. As we can tweak the settings of the instances run by OHD to match the capabilities of our routing algorithm, the outcomes are likely to be similar with sufficient runtime. However, due to practical reasons, we cannot let every routing instance run until an optimal solution is found. The larger an instance, the poorer the quality will be upon termination after a fixed amount of time. The difference between our outcome and the heuristic outcome for certain instances can be used as a measure to determine how reliable the conclusions of case III and case IV are for those instances. If we find a large difference, it is likely that the found outcomes do not reflect reality as much as we might like.

As the performance of the routing model is not our main interest, and the quality of these solutions are mostly relevant when evaluating the performance of the combined model, the simplest approach is to test the instances that will be used in later cases. So, for this case, we will run the same routing instances as in case III.

8.3.2. Case II: Charging Schedule + External VRP

The next experiments we would like to perform are to gauge the performance of the charging model when combined with routing solutions that are as realistic as possible. As our routing model is not needed for this case, we are not as limited by the instance size as we would be otherwise. We therefore propose two sets of experiments within this case. Firstly, we would like to test the performance of this model on the complete datasets. This means that we test instances as closely related to the original as possible. The only changes will be the fixed range, and the loosened vehicle-customer restrictions. Secondly, we would like to test the model on a reduced subset of instances. Ideally, we would like to compare the performance here to the performance of case III, but we believe that it is more valuable to test larger instances of 100 customers (generated in the same way as before), which cannot be solved efficiently by our routing algorithm.

To keep the comparison to the original dataset as faithful as possible, we will assume that the original split of electric and conventional vehicles are used. The goal is to compare policy (1) (see subsection 8.2.4) and policy (2) against policy (3) combined with the ability to charge vehicles in between. Additionally, we would also like to test the behavior of the purely electric fleet. Using policy (1) for this fleet is possibly infeasible due to the limited range of the EVs, so we will only use policy (2) to compare with policy (3) instead.

8.3.3. Case III: VRP + Charging Schedule

This case tests the both algorithms simultaneously, which allows the largest variability in settings for us to test. This means that it's extra important to carefully select the instances. As our routing model will need to solve these instances, an instance-size of 50 is used. We will compare the existing fleet, and the new fully electric fleet, as we did before. We will also make similar policy comparisons as in case II. We differ from case II however in the sense that we will test two versions of each of these instances: one in which we have applied post-processing to optimize the starting and ending times of the shifts, and one in which we did not. We also test model variation 3 for the mixed fleet, by assuming we have access to only a single slow charger. This is not as interesting for the fully electric fleet, as those vehicles can drive larger distances, requiring less vehicles overall.

8.3.4. Case IV: Vehicle and Charger Investments

In this case, we aim to illustrate the investment advice that our model can provide. This can be done with model variation 4. While not built for this purpose, the OHD engine can be used to determine which vehicles should be added to the fleet if the currently available vehicles are not sufficient. Using a subset of 50 customers, and assuming only two vehicles are freely available, while the remaining vehicles can be scheduled at their full purchase cost, should provide a sufficient baseline. As we are including the option to uplift as many chargers as possible, we are assuming that each vehicle can be fully charged before the start of the next shift and therefore make use of policy (3).

Unlike before, in this scenario we do assume we have access to a fleet that contains both the currently owned and the new vehicles. With the aim of phasing out the use of conventional vehicles, this fleet will consist of two of first type of EV that can be used at only a small cost, and a practically endless number of the new types of EV that can be used at purchasing cost.

9 Results

In this chapter, we discuss the outcomes for each of the tests for each of the four cases. Each case contains a group of instance types, which we will test for the four datasets. The performance of each of the instances will be determined not only by the objective value, but also by the indicators of how many vehicles are used, and how many kilometers are driven by EVs. For some instances we also provide the routing and charging schedules, to properly illustrate the behavior of the models.

9.1. Case I: VRP

This first case aims to test the performance of our routing model, when compared to the OHD engine that is currently used to solve the full-sized instances of our datasets. In this section we compare the results from the instances with 50 customers under policy (3), as these are the instances we are mainly interested in for case III. The instances studied in case II are too large to be run by our routing algorithm within a reasonable time frame, and the routing instances of case IV operate under slightly different assumptions and will be treated separately. This leaves us two sets of instances to compare, starting with the mixed fleet.

9.1.1. Mixed fleet

Table 9.1 summarizes the routing solutions generated by the OHD engine. Here, the total cost refers to the cost of each customer visit, each vehicle scheduled, and each kilometer driven by the scheduled vehicles. The electricity cost is simply the cost of the energy used by the vehicles driving these routes.

Table 9.1: Routing outcomes for the mixed fleet with 50 customers, generated by OHD.

Dataset	1	2	3	4
Total Cost	€$2,802.09$	€$3,814.15$	€$4,239.38$	€$2,005.11$
Number of ICEVs Used	3	3	3	2
Number of EVs Used	2	2	2	3
Percentage EVs	40%	40%	40%	60%
Total Kilometers Driven	543,758	541,200	648,242	533,074
Total Kilometers Driven by ICEVs	302,805	325,019	450,532	209,110
Total Kilometers Driven by EVs	240,953	216,181	197,710	323,964
Percentage of Kilometers Driven by EVs	44.3%	39.9%	30.5%	60.8%
Electricity Cost	€37.95	€34.05	€31.14	€51.02

Next, we provide the routing solutions from our exact solver in table 9.2. These had up to an hour and a half to solve, and terminated with a gap somewhere in between 3% and 36%.

Table 9.2: Routing outcomes for the mixed fleet with 50 customers, generated by our routing model.

Dataset	1	2	3	4
Total Cost	€353.76	€347.67	€$1,604.44$	€$1,436.42$
Number of ICEVs Used	0	0	1	1
Number of EVs Used	6	6	7	6
Percentage EVs	100%	100%	87.5%	85.7 %
Total Kilometers Driven	683,688	660,256	839,262	810,799
Total Kilometers Driven by ICEVs	0	0	152,485	129,318
Total Kilometers Driven by EVs	683,688	660,256	686,777	681,461
Percentage of Kilometers Driven by EVs	100%	100%	81.83%	84.05%
Electricity Cost	€107.68	€103.99	€108.17	€107.33

In table 9.3, the routing outcomes of our routing algorithm are compared against the solutions of the OHD engine:

Table 9.3: Comparison of the outcomes of our routing model and OHD.

Dataset	1	2	3	4
Cost difference	-87.4%	-90.9%	-62.1%	-28.4%
Difference in Total Vehicle Use	20%	20%	60%	40%
Difference in EV Use	300%	300%	350%	200%
Difference in Total Kilometers Driven	20.5%	18.0 %	22.8 %	34.2 %
Difference in EV Kilometers Driven	283.7 %	305.4 %	347.4 %	210.4 %

9.1.2. Fully Electric Fleet

Just as with the mixed fleet, table 9.4 summarizes the routing solutions generated by the OHD engine:

Table 9.4: Routing outcomes for the fully electric fleet with 50 customers for policy (3) generated by OHD.

Dataset	1	2	3	4
Total Cost	€464.93	€444.46	€474.25	€504.01
Number of $50kWh$ EVs	6	6	5	7
Number of $75kWh$ EVs	0	0	1	0
Total Kilometers Driven	695,884	617,156	712,515	696,200
Electricity Cost	€99.37	€88.13	€101.75	€99.42

Next, we provide the routing solutions from our exact solver in table 9.5. These instances again had up one and a half hour to solve, and terminated with a gap somewhere in between 1% and 3%.

Table 9.5: Routing outcomes for the fully electric fleet with 50 customers for policy (3) generated by our routing model.

Dataset	1	2	3	4
Total Cost	€417.10	€403.34	€480.46	€353.78
Number of $50kWh$ EVs	4	4	4	3
Number of $75kWh$ EVs	1	1	2	1
Total Kilometers Driven	642,677	590,141	717,168	549,17
Electricity Cost	€90.67	€83.33	€98.70	€77.10

In table 9.6, the routing outcomes of our routing algorithm are compared against the solutions of the OHD engine:

Table 9.6: Comparison of the outcomes of our routing model and OHD.

Dataset	1	2	3	4
Cost difference	- 10.3 %	- 9.2 %	1.3 %	-29.8 %
Difference in Vehicle Use	-16.7%	- 16.7%	0 %	-42.9%
Difference in Kilometers Driven	-7.6%	- 4.4 %	0.6%	-21.1 %

9.1.3. Analysis

Before analyzing the difference between these two solutions, it should be noted that the output from these solvers cannot be compared 1-to-1. The first reason for this is that our solver is exact, and the OHD engine solves the instances heuristically. On top of that, our solver is allowed 90 minutes to find a good solution, which was much longer than the OHD solver that spent only about 5 seconds on each of these instances. The OHD engine was also built to be able to solve large instances well, which is a very different task compared to solving small instances to optimality, and the same performance is not guaranteed, even if similar approaches are used. Another difference is that our exact routing model does not consider timing at all, apart from making sure the time-window constraints are satisfied. Post-processing is applied to make the schedules workable (i.e. move the activities closer together, add breaks and reduce the total duration time), which is done directly during the creation of the schedules by OHD. A final difference is that to keep the complexity of our routing model lower, we essentially remove the cost for adding vehicles to the model (which was added manually afterwards). This is not the case for the OHD engine. As a result, slightly different problems are solved.

Analyzing these results, we find that the quality of the routing solutions generated by our routing algorithm is quite high. The most remarkable difference between the two solvers is that our routing algorithm avoids using conventional vehicles at all cost, only scheduling a non-electric vehicle twice in the mixed fleet scenario, which was then only assigned a single customer. This is very different from the solutions given by OHD, which uses about as much EVs as ICEVs. The result is that the OHD solutions are more costly, but the vehicles drive less kilometers, on average $23,9\%$. The routes from our routing model however have a lower objective value, but the vehicles do drive much more kilometers overall, even though almost all of them are by an electric vehicle.

For the fully electric fleet, the differences between the two solvers are less stark. We see that our solver schedules the same number of vehicles, or one or two less. This also means a decrease in driven kilometers for all but one of the datasets, with an average of 8%.

Overall, we conclude that the quality of our routing solutions are sufficient to use for our analysis. In particular the routing solutions of the fully electric fleet resemble each other well enough. For the combined fleet, the differences between the two solvers are more apparent, so more care should be taken when drawing conclusions.

9.2. Case II: Charging Schedule + External VRP

The instances we are going to test for this case can be divided up into three categories. First, we want to test the full-sized instances, that are altered only such that the vehicle-customer restrictions are loosened, allowing electric vehicles to drive outside of the city center. Because of this, the EVs are now limited by the range of (half) their batteries. Next, we run the same tests for an instance size of 100 customers. These instances are not solved by our own routing algorithm due to its size, but they do illustrate the performance of the routing model better than the full-sized, or the smaller instances. Finally, we are going to do these same tests when considering a fleet of only EVs. When we have access to a mixed fleet, we compare policy (1) and (3), and policy (2) and policy (3). When the fleet we are using is fully electric, we only make the latter comparison. For policy (3) to be feasible, we need to make sure that all vehicles are sufficiently charged. This is where the charging schedule comes in: if we can find charging schedules such that all vehicles are charged as much as they need to, then we are sure that all the vehicles are able to finish their routes.

In total, there are three charging instances: charging during the day in between instance 1 and 2, and 3 and 4, and charging overnight between instance 2 and 3. For the full-sized instances we assume to have access to 22 $22kW$ chargers, and 4 $50kW$ chargers. For the smaller instances, we assume to have access to 4 $22kW$ chargers, and 2 $50kW$ chargers instead. This results in a total of respectively 26 and 6 chargers, which is roughly half of the scheduled EVs for each instance type. Each set of instances is tested under two assumptions:

1. Charge only as much as necessary, or
2. Charge as much as possible

This first assumption finds the bare minimum amount of charging needed to be able to successfully complete all routes. In the second assumption we assume that each vehicle is to be charged as much as possible using the available chargers. This might not be the most cost-efficient method of charging when considering individual instances, but it does provide the most leniency for the drivers and flexibility for possible last-minute schedule changes. These costs are also closer to the total energy costs attached to using the EVs during a shift. The charging costs under the first assumption only cover the bare minimum to perform an extra shift during the day, and are hence not representative of the actual costs associated with charging. These costs are also calculated for each of the instances. We mention in chapter 8 that we had very little information about the used vehicles, so the expended energy per kilometer driven is merely an approximation. This nevertheless gives an idea of the energy costs.

9.2.1. Full-sized datasets

When using the full datasets, we only consider the first and third policy. We start by comparing the routing solutions from both policies without the interference of the charging schedule. These results are given for policy (1) in table 9.7 and for policy (3) in table 9.8. It should be noted that for these two tables, the total cost refers to a slightly different objective function than before. As argued in section 8.2, the hourly driver rates were omitted from the reduced instances. Here, these are included.

Table 9.7: Routing outcomes for the full-sized dataset for policy (1) generated by OHD.

Dataset	1	2	3	4
Number of ICEVs	120	109	107	125
Number of EVs	54	54	54	54
Total Cost	€$210,630.94$	€$196,985.04$	€$181,711.38$	€$193,080.35$
Number of ICEVs Used	87	76	74	91
Number of EVs Used	54	54	54	54
Percentage EVs	38.3%	41.5%	42.2%	37.2%
Total Kilometers Driven	7,205,817	7,071,046	6,480,972	6,601,080
Total Kilometers Driven by ICEVs	4,258,648	4,093,834	3,699,255	3,636,649
Total Kilometers Driven by EVs	2,947,169	2,977,212	2,781,717	2,964,431
Percentage of Kilometers Driven by EVs	40.1%	42.1%	42.9%	44.9%
Electricity Cost	€462.94	€467.66	€436.95	€465.65

Table 9.8: Routing outcomes for the full-sized dataset for policy (3) generated by OHD.

Dataset (Full-sized)	1	2	3	4
Number of ICEVs	120	109	107	125
Number of EVs	54	54	54	54
Total Cost	€$198,706.13$	€$181,567.60$	€$174,894.32$	€$190,776.03$
Number of ICEVs Used	63	70	70	90
Number of EVs Used	54	54	54	54
Percentage EVs	45.1%	43.5%	43.5%	37.5%
Total Kilometers Driven	7,383,359	7,465,935	6,959,898	7,254,713
Total Kilometers Driven by ICEVs	3,802,684	3,514,501	3,482,845	3,363,816
Total Kilometers Driven by EVs	3,580,675	3,951,434	3,477,053	3,890,897
Percentage of Kilometers Driven by EVs	48.5%	52.3%	50.0%	53.6%
Electricity Cost	€562.45	€620.69	€546.17	€611.18

Table 9.9: Comparison of the outcomes for policy (1) and policy (3) for the full-sized dataset.

Dataset	1	2	3	4
Cost Reduction	€$11,924.81$	€$15,417.44$	€$6,817.06$	€$2,304.32$
Percentual Cost Reduction	5.7%	7.8%	3.8%	1.2%
Reduction in ICEVs Used	24	6	4	1
Percentual ICEV Use Reduction	27.6%	7.9%	5.4%	1.1%
Increase in Kilometers Driven by EVs	633,027	974,222	695,336	926,466
Decrease in Kilometers Driven by ICEVs	455,964	579,333	216,410	272,833
Percentual Increase of Kilometers Driven by EVs	21.5%	32.7%	25.0%	31.2%
Percentual Decrease of Kilometers Driven by ICEVs	10.7%	14.2%	5.9%	7.5%
Increase in Total Kilometers Driven	117,542	394,889	478,926	653,633
Percentage Point Increase of EV Kilometers	8.4%	10.1%	7.1%	8.7%
Percentual Increase of Kilometers Driven	2.5%	5.6%	7.3%	9.9%

Comparing policy (1) with policy (3), we see some great improvements. On average (over all datasets) we see that moving from policy (1) to (3) results in a cost reduction of 4.6%. This cost reduction is explained mostly by an increase of kilometers driven by EVs of on average 27.6% (against an average 6.3% increase in total driven kilometers), in addition to an average 10.5% decrease of non-electric vehicles used. This lead to an average decrease of kilometers driven by non-electric vehicles of 9.6%. We also see that on average, the percentage of kilometers driven by EVs grows from 42.5% to 51.1%, an increase of 20.2%. We may additionally note that the average driven distance for this policy is 69 kilometers, compared to 54 kilometers for policy (1). This is a significant increase, but as these vehicles are allowed to drive up to 133 kilometers during a single trip, it might be possible to see even larger improvements by scheduling the vehicles even more.

In the next two tables, a summary of the charging schedule is provided for each of the two assumptions mentioned above. Table 9.10 covers assumption 1, while table 9.11 covers assumption 2.

Table 9.10: Charging outcomes for the routing pairs of the full-sized dataset under assumption 1.

Charging Instance	1	2	3
Chargers Used	0	2	4
Slow Chargers Used	0	0	1
Fast Chargers Used	0	2	3
Charger Operations	0	6	8
Total Electricity Used	$0kWh$	$15.50kWh$	$22.52kWh$
Total Electricity Price	€0	€6.51	€9.46

Table 9.11: Charging outcomes for the routing pairs of the full-sized dataset under assumption 2.

Charging Instance	1	2	3
Chargers Used	23	22	26
Slow Chargers Used	19	18	22
Fast Chargers Used	4	4	4
Charger Operations	68	98	76
Total Electricity Used	$1,019.17kWh$	$1,453.09kWh$	$1,104.07kWh$
Total Electricity Price	€428.05	€610.30	€463.71

We see that under assumption 1, we need very little, or even no chargers to complete the two shifts. That is because the vehicles in most of these datasets drive short distances (the average over all datasets for this policy is 54 kilometers), and are able to drive two shifts using a full battery. At the end of the day, all vehicle batteries will be mostly depleted. Assumption 2 does charge the vehicles as much as possible, and as a result allows the vehicles to start the second shift with much more energy than before. When comparing the cost of electricity that is necessary to drive the vehicles that are routed, and the cost of electricity that is spent during the charging phase, we see that vehicles after charging are left with most of the energy that they started with, on average. To be exact, we respectively charge 76.1%, 98.3% and 84.9% of the energy used in the shift right before its starting moment. Ideally, we want the second charging instance to charge the vehicles fully. Here, despite the fact that not all chargers are occupied, this does not happen. We expect that this happened as a consequence of terminating the model prematurely, and in realistic instances, the vehicles should be able to be fully charged overnight with the current amount of chargers.

9.2.2. Mixed Fleet, 100 Customers

Now, we move on to reduced datasets. These datasets contain 100 customers, and miss some time-related factors, as described in section 8.2. We do however use the same fleet as before. We once again summarize the outcomes of policy (1) in table 9.12 and policy (3) in table 9.13 below, followed by a comparison between these two policies in addition to policies (2) and (3) in table 9.14. Policy (2) closely resembles policy (3), with the exception that we are forced to use vehicles that have not been used previously. This means that for this policy there is no separate set of routing solutions.

Table 9.12: Routing outcomes for the mixed fleet with 100 customers for policy (1) generated by OHD.

Dataset (100 mixed)	1	2	3	4
Total Cost	€$12,140.97$	€$11,765.59$	€$11,853.15$	€$11.853,30$
Number of ICEVs Used	8	8	8	8
Number of EVs Used	0	1	1	0
Percentage EVs	0%	11.1%	11.1%	0%
Total Kilometers Driven	902,005	859,889	887,892	855,690
Total Kilometers Driven by ICEVs	902,005	793,784	825,462	855,690
Total Kilometers Driven by EVs	0	66,105	62,430	0
Percentage of Kilometers Driven by EVs	0%	7.7%	7.0%	0%
Electricity Cost	€0	€10.41	€9.83	€0

Table 9.13: Routing outcomes for the mixed fleet with 100 customers for policy (3) generated by OHD.

Dataset (100 mixed)	1	2	3	4
Total Cost	€$5,839.97$	€$5,389.69$	€$2,226.18$	€$5,455.40$
Number of ICEVs Used	3	3	2	2
Number of EVs Used	5	5	6	4
Percentage EVs	62.5%	62.5%	75%	66.7%
Total Kilometers Driven	868,474	827,241	872,599	780,164
Total Kilometers Driven by ICEVs	336,582	298,321	210,505	352,395
Total Kilometers Driven by EVs	531,892	528,920	662,094	427,769
Percentage of Kilometers Driven by EVs	61.2%	63.9%	75.9%	54.8%
Electricity Cost	€83.77	€83.30	€104.28	€67.37

Table 9.14: Comparison of the outcomes for policy (1) and policy (3) for the mixed fleet with 100 customers.

Dataset	1	2	3	4
Cost Reduction	€$6,301.00$	€$6,375.90$	€$9,626.97$	€$6.397,9$
Percentual Cost Reduction	51.9%	54.2%	81.2%	54.0%
Increase in EVs Used	5	4	5	3
Decrease in ICEVs Used	5	4	6	6
Percentual EV Use Increase	-	400%	500%	-
Percentual ICEV Use Decrease	65.5%	65.5%	75%	75%
Increase in Kilometers Driven by EVs	531,892	462,815	599,664	427,769
Decrease in Kilometers Driven by ICEVs	565,423	495,463	614,957	503,295
Percentual Increase of Kilometers Driven by EVs	-	800,1%	1060,5%	-
Percentual Decrease of Kilometers Driven by ICEVs	62.7%	62.4%	74.5%	58.8%

The difference between policy (1) and policy (3) is even more pronounced for these instances, compared to using the full dataset. On average, we find a cost reduction of $60,3$%, made possible by the introduction of up to five EVs in each dataset. Seeing how under policy (1) none, or only a single EV is used, this is a great improvement. We find that on average, this implies that the number of kilometers driven by ICEVs decreases by 64.6%. When policy (2) is used, this difference is not expressed as a difference in cost or driven kilometers, but by the additional vehicles that are needed to drive the routes as planned. We find that five extra vehicles are needed, up to 100% more EVs than initially scheduled.

It it also possible to compare the baseline results of policy (2) and policy (3), which is done in table 9.15. We study these instances case by case, and assume that we have access to the vehicles that are scheduled in the first shift. To draw conclusions about the entire scenario, one can simply take the maximal set for each vehicle type over all instance pairs.

Table 9.15: Comparison of the outcomes for policy (2) and policy (3) for the mixed fleet with 100 customers.

Charging Instance	1	2	3
ICEVs Used Policy 3	3	3	2
EVs Used Policy 3	5	5	6
ICEVs Used Policy 2	3	3	2
EVs Used Policy 2	10	10	9
Extra EVs Needed	5	5	4
Percentual EV Increase	100%	100%	66,7%

It should be noted that the second pair of instances, namely the late shift on the first day, and the early shift on the second day, has a lot of time in between, which is typically used to charge the vehicles. In this table we consider these pairs of instances individually, to see how many vehicles need to be added if we do not get the opportunity to charge between the shifts. If we consider the bigger picture, that is, allowing the vehicles to fully charge overnight, this instance pair does not require any additional vehicles.

Now we involve the charging schedules, and provide a summary for both assumption 1 in table 9.16 as well as assumption 2 in table 9.17:

Table 9.16: Charging outcomes for the routing pairs of the the mixed fleet with 100 customers under assumption 1.

Charging Instance	1	2	3
Chargers Used	4	3	4
Slow Chargers Used	2	1	2
Fast Chargers Used	2	2	2
Charger Operations	8	10	8
Total Electricity Used	$115.69kWh$	$137.33kWh$	$115.21kWh$
Total Electricity Price	€48.59	€57.68	€48.39

Table 9.17: Charging outcomes for the routing pairs of the mixed fleet with 100 customers under assumption 2.

Charging Instance	1	2	3
Chargers Used	5	1	5
Slow Chargers Used	3	0	3
Fast Chargers Used	2	1	2
Charger Operations	16	10	14
Total Electricity Used	$199,45kWh$	$198.35kWh$	$248.29kWh$
Total Electricity Price	€83.77	€83.30	€104.28

Comparing the cost of electricity that is needed to drive the scheduled vehicles, and the cost of charged electricity during the charging phase, we see that vehicles can be charged fully before the start of their next shift. When this is not the goal, we see that the vehicles still need to be charged on average past their halfway point. Unlike the full-sized instances that hardly require any charging between shifts, this shows us that policy (3) does depend quite heavily on the opportunity to charge vehicles during the day, when the option to schedule additional EVs exists.

9.2.3. Fully Electric Fleet, 100 Customers

Now, we run the same instances, except this time we make use of a fully electric fleet. As policy (1) is not feasible for this fleet, we assume that the base policy would be policy (2). Table 9.18 summarizes the results of policy (3), the comparison between the two policies is given in table 9.19.

Table 9.18: Routing outcomes for the fully electric fleet with 100 customers for policy (3) generated by OHD.

Dataset	1	2	3	4
Total Cost	€723.98	€714.81	€694.29	€701.34
Number of $50kWh$ EVs	8	8	8	8
Number of $75kWh$ EVs	1	1	1	1
Total Kilometers Driven	1,030,706	995,437	916,494	943,615
Electricity Cost	€145.96	€141.00	€129.54	€133.45

To compare the situation with the charging schedule against the situation without the opportunity to charge, we assume that unless a shift can be driven by one of the vehicles that was used previously, an extra vehicle is used. We can only make this comparison when looking at pairs of instances. Because each instance used the same set of vehicles, these vehicles are available for each instance pair.

Table 9.19: Comparison of the outcomes for policy (2) and policy (3) for the fully electric fleet with 100 customers.

Charging Instance	1	2	3
$50kWh$ EVs Used Policy 3	8	8	8
$75kWh$ EVs Used Policy 3	1	1	1
$50kWh$ EVs Used Policy 2	12	12	12
$75kWh$ EVs Used Policy 2	2	2	2
Extra EVs Needed	5	5	5
Percentual EV Increase	55.5%	55.5%	55.5%

As mentioned in section 9.2.2, this second charging instance is typically used to charge all the vehicles, as the vehicles are stationary the entire night. These values simply show what would happen if these two instances are scheduled closely together. If we consider these instances sequentially, we do not need to add additional vehicles for this instance pair. Considering that each instance pair requires a total of fourteen EVs, this does not impact the minimal fleet size needed.

We are slightly limited in the extent to which we can draw conclusions when compared to the previous sets of instances, as we are unable to compare the outcomes from policy (3) against policy (1). We do see from the outcome of policy (2) that not allowing the vehicles to charge will require a set of EVs that is 55.5% larger than otherwise needed, increasing the fleet size from nine to fourteen to visit all the customers in these instances.

The following tables summarizes the charging outcomes for each of the two assumptions. Table 9.20 covers the results using assumption 1, the results for assumption 2 are given in table 9.21.

Table 9.20: Charging outcomes for the routing pairs of the fully electric fleet with 100 customers under assumption 1.

Charging Instance	1	2	3
Chargers Used	2	2	3
Slow Chargers Used	0	1	1
Fast Chargers Used	2	1	2
Charger Operations	8	6	6
Total Electricity Used	$143.45kWh$	$112.0kWh$	$102.33kWh$
Total Electricity Price	€60.25	€47.04	€42.98

Table 9.21: Charging outcomes for the routing pairs of the fully electric fleet with 100 customers under assumption 2.

Charging Instance	1	2	3
Chargers Used	6	6	6
Slow Chargers Used	4	4	4
Fast Chargers Used	2	2	2
Charger Operations	26	24	22
Total Electricity Used	$357.50kWh$	$339.19kWh$	$270.07kWh$
Total Electricity Price	€145.96	€141.00	€113.43

From these tables, we can see that for two of the three instance pairs, the used vehicles are able to be fully charged before their next shift. The final pair is able to charge 85.0% of the energy used in the previous shift, allowing the vehicles to still leave for the final routing instance mostly charged. If we are only interested in charging the vehicles enough so that they can drive their next shift, we see that we only require up to half as many chargers as otherwise: only two or three chargers are needed for this fleet consisting of nine EVs to drive another shift that same day.

9.3. Case III: VRP + Charging Schedule

When using our exact routing model, we cannot solve large instances within reasonable time. For that reason, we only solve instances of 50 customers, for both types of fleets. These are also the routing instances that are tested in case I. The set of available chargers is determined in the same way as before, resulting in the use of two or four $22kW$ chargers, depending on the fleet type, and two $50kW$ chargers. In case II, we assume that vehicles start their shift with full batteries. This is done because those routing solutions are calculated independently: no information about the shifts earlier that day is known when the routes for the shift later that day are calculated. For that reason, we treated each of the charging instances as independent and hence did not assume that the battery levels of the previous day carried onto the next. Now that we are calculating the routes ourselves, we see them as consecutive events. For most of the tests performed here, we still assume that the vehicles have access to a full battery when the routes are calculated, and we hope that the vehicles can be charged sufficiently before they need to drive their shifts.

For the mixed fleet, we perform an additional test. In this test, we assume that we are not able to fully charge our vehicles. To still find feasible routes, we will estimate the battery capacity of the vehicles before calculating the routes for each instance (apart from the first instance, which we still assume starts with a fully charged fleet). To properly illustrate this functionality, we will assume that a much smaller set of chargers is available. This means that we know that not all vehicles can be charged sufficiently, so as a consequence the routing algorithm will need to adapt.

9.3.1. Mixed fleet

In case I we discuss how the solutions of our routing model and the OHD engine differ for these datasets, but we have not yet analyzed how our outcomes line up with the solutions calculated under policy (1).

Here, we present a summary of our routing solutions when policy (1) is used in table 9.22, and follow that by the table presented earlier that contains our outcomes for policy (3), repeated in table 9.23. These two tables are then compared in the same way as in case II. Those results are given in table 9.24.

Table 9.22: Routing outcomes for the mixed fleet with 50 customers for policy (1) generated by our routing model.

Dataset	1	2	3	4
Total Cost	€$2,616.93$	€$3,667.97$	€$3,504.74$	€$3,059.95$
Number of ICEVs Used	1	3	2	2
Number of EVs Used	10	7	7	6
Percentage EVs	90.9%	70.0%	77.7%	75.0%
Total Kilometers Driven	866,604	795,628	712,223	639,694
Total Kilometers Driven by ICEVs	218,426	375,595	357,864	313,131
Total Kilometers Driven by EVs	648,178	420,033	354,359	326,563
Percentage of Kilometers Driven by EVs	74.8%	52.8%	49.7%	51.0%
Electricity Cost	€102.09	€66.15	€55.81	€51.43

Table 9.23: Routing outcomes for the mixed fleet with 50 customers for policy (3) generated by our routing model.

Dataset	1	2	3	4
Total Cost	€353.76	€347.67	€$1,604.44$	€$1,436.42$
Number of ICEVs Used	0	0	1	1
Number of EVs Used	6	6	7	6
Percentage EVs	100%	100%	87.5%	85.7 %
Total Kilometers Driven	683,688	660,256	839,262	810,799
Total Kilometers Driven by ICEVs	0	0	152,485	129,318
Total Kilometers Driven by EVs	683,688	660,256	686,777	681,461
Percentage of Kilometers Driven by EVs	100%	100%	81.83%	84.05%
Electricity Cost	€107.68	€103.99	€108.17	€107.33

Table 9.24: Comparison of the outcomes for policy (1) and policy (3) for the mixed fleet with 50 customers.

Dataset	1	2	3	4
Cost Reduction	€$2,263.17$	€$3,304.11$	€$1,900.30$	€$1,623.53$
Percentual Cost Reduction	86.5%	90.1%	54.2%	53.1%
Reduction in EVs Used	4	1	0	0
Percentual EV Use Reduction	40%	14.3%	0%	0%
Reduction in ICEVs Used	1	3	1	1
Percentual ICEV Use Reduction	100%	100%	50%	50%
Increase in Kilometers Driven by EVs	35,510	240,223	250,588	270,848
Decrease in Kilometers Driven by ICEVs	218,426	375,595	205,379	183,813
Percentual Increase of Kilometers Driven by EVs	5.2%	36.4%	36.5%	39.7%
Percentual Decrease of Kilometers Driven by ICEVs	100%	100%	57.4%	58.7%
Reduction of Total Kilometers Driven	182,916	135,372	25,446	-41,767
Percentual Reduction of Total Kilometers Driven	26.8%	20.5%	3.7%	-6.1%

Unlike the two comparisons between these two policies in case II in section 9.2.1 and section 9.2.2, we do not unambiguously find an increase or a decrease of the total number of kilometers driven. For most instances this distance decreases, by up to 26.8%, but for the final instance we find that the decrease in ICEVs used lead to more inefficient routes. This comparison also stands out from the others as the number of EVs used decreases for half of the instances, and didn't change for the others. From this we conclude that more so than the OHD engine, our routing model tries to schedule as many EVs as possible, no matter the range. When the range increases, less EVs might be necessary. This follows from seeing that in one instance the number of EVs scheduled decreases by 40%, while for two others it does not decrease at all. Despite the lower number of vehicles used, they do drive more kilometers on average: we see an average increase of 29.5% of kilometers driven by EVs. What is not unexpected about these results is the fact that for all the instances, the number of ICEVs used decreases, for some instances even to exclusively scheduling EVs. As a result, we find a 79.0% decrease in kilometers driven by non-EVs. The average cost reduction of 71% is also in line with our conclusions from section 9.2.2.

Sufficient Charging Capacity
As mentioned before, unlike in case II, only the first instance is assumed to have a fully charged fleet. To distribute the charging load as evenly as possible, vehicles are be charged maximally during their charging intervals. This avoids situations in which vehicles are minimally charged during the first charging session, and need to be charged overnight from their completely drained state. Using these assumptions, the charging outcomes can be summarized by the tables below. Table 9.25 contains the results for the charging problem when no post-processing was applied. In table 9.26, the results of the same problem are given, except this time post-processing was applied, increasing the length of the charging windows.

Table 9.25: Charging outcomes for the routing pairs of the mixed fleet with 50 customers, without post-processing.

Charging Instance	1	2	3
Chargers Used	3	2	4
Slow Chargers Used	2	2	2
Fast Chargers Used	2	2	2
Charger Operations	10	12	12
Vehicles Added	1	1	0
Total Electricity Used	$183.33kWh$	$279.00kWh$	$203kWh$
Total Electricity Price	€77.00	€117.18	€85.26

Table 9.26: Charging outcomes for the routing pairs of the mixed fleet with 50 customers, with post-processing.

Charging Instance	1	2	3
Chargers Used	2	4	3
Slow Chargers Used	0	2	1
Fast Chargers Used	2	2	2
Charger Operations	12	12	10
Vehicles Added	0	2	0
Total Electricity Used	$256.28kWh$	$247.69kWh$	$216.24kWh$
Total Electricity Price	€107.64	€104.03	€90.82

Comparing the energy used in the the charging schedule ran after the application of post-processing, we see that for the first two instances 100% of the energy used in the previous shift is recuperated. The energy costs for these separate instances aren't identical; the first charging instance is short €0.04, which is added in the next charging opportunity. During last the charging interval we are not able to charge all instances fully before the final shifts start, instead charging 83.5% of the energy used in the previous shifts.

When post-processing isn't applied, resulting in shorter charging intervals, we find that the first and last charging interval respectively charged 71.5% and 78.8% of the energy used in the previous shifts. This is partly because there was insufficient opportunity to charge, but also in part because one of the shifts is completed by a conventional vehicle, that cannot be charged. The middle, overnight, charging interval does fully charge the vehicles before the morning shift. The difference between the sum of the energy costs of the first two charging instances, and the sum of the energy costs of the first two routing instances, is due to the conventional vehicle that drives a shift originally scheduled for an electrical vehicle.

Another remarkable thing is that while we do need to add (at least) two vehicles at some point during this sequences of shifts, as the third set of routes contains eight shifts, two more than the six shifts of the initial set of routes, this is done differently in both calculations. In the first table we see that a vehicle is added during the first charging moment, as the shifts are planned too tightly. This is solved after applying post-processing, in which case the extra vehicles are only added right before the number of shifts increases.

Insufficient Charging Capacity

To properly test model variation 3, we need to restrict the set of available chargers. Otherwise, the battery estimates made would hardly affect the routing outcomes. In section 9.4 we will see that two $22kW$ chargers are sufficient to charge the fleet used in case IV. We find that this is also sufficient for the mixed fleet in case III, and therefore choose to only allow a single $22kW$ charger for this experiment. As a baseline, we use the routes found by our model after post-processing was applied. If routes cannot be completed by the vehicles that did not get to charge sufficiently, other vehicles will be scheduled. The total number of vehicles scheduled in the different scenarios are given in table 9.27. The first two rows give the minimal number of vehicles needed, according to the routes summarized by table 9.23. The next two rows contain the number of vehicles that were scheduled by our charging model. Note that not all of these vehicles are used every charging instance, they can also be left on standby, and be used to drive later routes.

The main scenario only allows the routing model access to the vehicles used by the original routing instances, in addition to several ICEVs. Previously, we allowed the model access to the entire fleet of over 160 vehicles, as only as many vehicles would be scheduled as needed. Here however, access to this many vehicles means that the model would simply schedule vehicles that do not have a partially drained battery. To avoid that, we need to limit the set of vehicles that the model may schedule. Limiting this set too strictly will however result in infeasible solutions, which cannot be used in practice. The compromise we used was to allow a fleet of seven EVs, expanded by several ICEVs. The idea was that these ICEVs would only be used if none of the EVs had sufficient range (left).

Table 9.27: Number of vehicles scheduled per charging instance for different modeling scenarios

Charging Instance	1	2	3
Minimal number of EVs needed	6	7	6
Minimal number of ICEVs needed	0	1	1
Number of EVs scheduled for the Baseline Scenario	8	8	8
Number of ICEVs scheduled for the Baseline Scenario	1	1	2
Number of EVs scheduled for Main Scenario	3	4	4
Number of ICEVs scheduled for Main Scenario	3	3	3

We see that the number of vehicles scheduled in the baseline scenario is much higher than needed. Overall, one ICEV and seven EVs are needed to complete all four routing instances, but due to the limited charger availability, eight EVs and two ICEVs were scheduled, resulting in a total of ten used vehicles. The main scenario, that makes use of the battery estimates, takes a different approach and schedules three ICEVs, and four EVs in total. Now, only seven vehicles are necessary. It is not obvious to conclude which of these schedules is better, as three of these seven vehicles are ICEVs. The reason for this behavior is that in the first charging instance, three ICEVs were scheduled, as the only vehicles that this model were allowed to add were vehicles used previously. These vehicles were then repurposed later on, driving shifts that could have been driven by EVs, without the charging demand. We conclude that while model variation 3 did fulfill its purpose by finding a feasible schedule when insufficient charging was available, this was not necessarily the outcome were were aiming for. Under different assumptions that would avoid the immediate scheduling of three ICEVs, this model variation can still be very promising, but more experiments are needed to confirm this.

9.3.2. Fully Electric Fleet

Just like in case II, for the fully electric fleet we compare the results formed under policy (3) with those that follow from using policy (2). These results are presented in table 9.28. So, we do not present new routing results, and directly start analyzing how many extra vehicles would be needed were this policy used.

Table 9.28: Comparison of the outcomes for policy (2) and policy (3) for the fully electric fleet with 50 customers.

Charging Instance	1	2	3
$50kWh$ EVs Used Policy 3	4	4	3
$75kWh$ EVs Used Policy 3	1	2	1
$50kWh$ EVs Used Policy 2	8	7	6
$75kWh$ EVs Used Policy 2	2	2	2
Extra EVs Needed	5	4	4
Percentual EV Increase	100%	66.7%	100%

We see that for two out of the three charging instances, none of the vehicles are able to drive a second shift. The third instance is able to re-use two out of the five vehicles that are scheduled the previous shifts. As this set of shifts finishes during the evening, they typically do get charged before the next shift starts. For our conclusion this does not matter however, as we need a total of ten vehicles to successfully schedule all the shifts in these instances.

The outcomes of the charging model, with (table 9.29) and without post-processing (table 9.30) are given below.

Table 9.29: Charging outcomes for the routing pairs of the fully electric fleet with 50 customers, without post-processing.

Charging Instance	1	2	3
Chargers Used	3	4	3
Slow Chargers Used	1	2	0
Fast Chargers Used	2	2	2
Charger Operations	10	12	12
Extra Vehicles Added	1	0	0
Total Electricity Used	$162.83kWh$	$252.37kWh$	$225.73kWh$
Total Electricity Price	€68.39	€106.00	€94.81

Table 9.30: Charging outcomes for the routing pairs of the fully electric fleet with 50 customers, with post-processing.

Charging Instance	1	2	3
Chargers Used	3	4	3
Slow Chargers Used	1	2	1
Fast Chargers Used	2	2	2
Charger Operations	10	10	10
Extra Vehicles Added	0	1	0
Total Electricity Used	$215.60kWh$	$194.33kWh$	$235.92kWh$
Total Electricity Price	€90.55	€81.62	€99.09

To compare these two tables, it is important to keep in mind that the first charging instance schedules the same amount of shifts as available vehicles, the second charging instance has an extra shift to assign (and hence requires an extra vehicle to be added) and the third charging instance has less shifts to assign to the same amount of vehicles. This means that after post-processing, the quality of the charging solution is quite good: it only adds an extra vehicle if there are too many shifts (charging instance 2), and is able to charge the vehicles more when there are vehicles that will not be assigned a shift (charging instance 3). In the schedules where post-processing is not applied, we see that we add a vehicle to complete the second set of routes, when it was not strictly necessary. The reason for this is because the shifts are planned too tightly, and would have overlapped if not for the use of an extra vehicle.

Studying the amount of energy left in the batteries after having charged, we find that when we do use post-processing, there is sufficient time to charge, and we can fully charge the vehicles before every charging instance. It may be noted that the energy costs as presented in the tables do not match up exactly with the projected energy costs in section 9.1.2. The reason for this is that the two vehicle types have a slightly different energy consumption, and shifts might not be completed by the same type of vehicle that it was scheduled as initially. In the version where post-processing is not applied, the length of the charging interval does impact the feasibility of full vehicles upon departure. The vehicles can be charged fully overnight, but during the first and final charging opportunity this is not possible. We find that respectively 75.4% and 96.0% of the energy used during the previous set of routes was charged before the next shift started. The difference between these two values is explained by the fact that the third routing instance contained an additional vehicle that does not need to be scheduled, and could hence be charged for a longer amount of time. Because the vehicles can be charged fully overnight, such a schedule is still sustainable, but the added tightness does make it more difficult to be able to assign vehicles to shifts in a desirable manner, showing the main benefit of slightly bigger charging windows.

9.4. Case IV: Vehicle and Charger Investments

For this final case we test a single instance type. These instances contain 50 customers, and have access to two of the currently owned EVs, and the ability to purchase vehicles of the new type with both the smaller or the larger battery capacity. We also assume two $22kW$ chargers are available, with the option to purchase either $22kW$, $50kW$ or $150kW$ chargers. First we compare the outcomes of the routing algorithms, and afterwards we discuss the conclusions of the charging model.

9.4.1. Routing

Both the OHD engine and the exact routing model can be used to draw conclusions about what vehicles to purchase. This can even be done using the standard versions of both outcomes, when an abundance of vehicles was available. In these outcomes however, the price of the vehicles has had a much smaller impact on the choice of scheduled vehicles. When optimal solutions are not found, it is uncertain whether the set of used vehicles is indeed optimal cost-wise, or not. For that reason, increasing the incentive of the model to use cheaper vehicles can yield interesting results. We start with the solutions from the OHD solver, which are provided in table 9.31:

Table 9.31: Routing outcomes for the upliftable with 50 customers, generated by OHD.

Dataset	1	2	3	4
Total Cost	€$162,061.00$	€$162,053.69$	€$169,519.50$	€$202,519.18$
Number of $50kWh$ EVs Purchased	4	4	3	5
Number of $75kWh$ EVs Purchased	0	0	1	0
Total Kilometers Driven	650,001	621,880	682,692	681,471

The outcomes of our own routing algorithm are summarized below in table 9.32.

Table 9.32: Routing outcomes for the upliftable with 50 customers, generated by our routing model.

Dataset	1	2	3	4
Total Cost	€$169,506.71$	€$209,951.12$	€$169,545.03$	€$129,554.16$
Number of $50kWh$ EVs Purchased	3	4	3	2
Number of $75kWh$ EVs Purchased	1	1	1	1
Total Kilometers Driven	656,423	631,625	718,034	579,084

We see that unlike the fully electric fleet in case III in section 9.3.2, the OHD routing model significantly outperforms our exact model. We even see that this variation of the model outperforms the version of case I in section 9.1.2, driving less kilometers. The reason that our model performs worse, compared to both the solutions from case I and the OHD solutions, is because this version is more complex than the standard version. This was not accounted for in the allowed running time, as both models were terminated after 90 minutes. For all but one of the datasets, this version of the model used more EVs than the version in case I, allowing us only to conclude that that version of the model performs better at the task of determining the cheapest set of vehicles to purchase in order to visit all the customers within their time windows.

9.4.2. Charging

As is the case with routing, conclusions about which chargers to purchase can be drawn from the previous iterations of the charging model, in particular when not all available chargers were used. This is even more inaccurate than with vehicles however, as the differences between two vehicle batteries are much smaller in both cost as well as modeling benefit compared to the difference between two types of chargers.

Unlike the previous cases, we are trying to find a set of chargers that is as cheap as possible, and can charge all vehicles such that they can be used repeatedly. We see in section 9.2 that to charge vehicles between shifts only enough to be able to drive the next shift, we don't need very many chargers. We found that for 100 customers, one to four chargers are sufficient. That however assumes that the vehicles started the day with a full battery, and the vehicles do need to be charged up to that point overnight as well. In this section we will analyze the set of chargers that is necessary to complete both of these tasks: charging vehicles for their next shift, and charging them fully overnight.

A quick calculation shows that to charge six vehicles with $50kWh$ batteries fully, seven hours and two $22kW$ chargers are sufficient. If the total charging load is higher, when an extra vehicles is added (or two vehicles have the larger $75kW$ battery) resulting in a $350kW$ charging load, those same two chargers need a total of eight hours, which is available in all typical situations. This is the same conclusion that follows if you allow the model to run between any of the evening shifts, and any of the morning shifts. Figure 9.1 shows a possible charging schedule between the shifts from the second dataset, starting in the afternoon, and the shifts from the first dataset, starting in the morning. While this implies that in theory, charging overnight can be done very efficiently, employers might in practice prefer schedules that do not require several charger connections and disconnections in the middle of the night. This means that in its current state, the charging model is not a useful tool for determining the right set of chargers needed for overnight charging. To do this properly, more information is needed about the logistical restrictions regarding overnight charging.

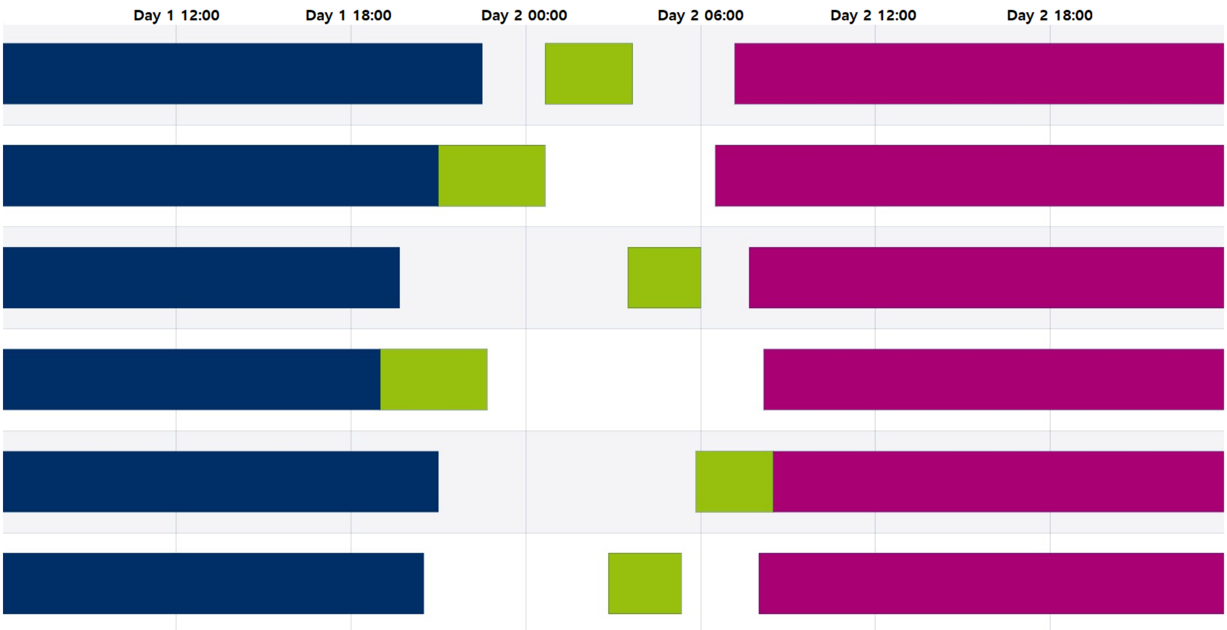

Figure 9.1: Charging Schedule of time between the second and the first set of routes.

It is possible that more, or higher capacity, chargers might be needed to sufficiently charge the vehicles during the day. For the shifts created by the OHD engine however, this turned out to not be necessary. For both sets of early and late instances, two $22kW$ chargers were sufficient to be able to use the vehicles for the next shift. We also considered the two other pairings of early and late shifts (instance 4 and 1, and instance 3 and 2), and found the same result. Of this final example, we provide the charging schedule in figure 9.2 to illustrate the behavior. We see that the charging windows are smaller, but as the charging demand is also lower, we are still able to charge all batteries fully.

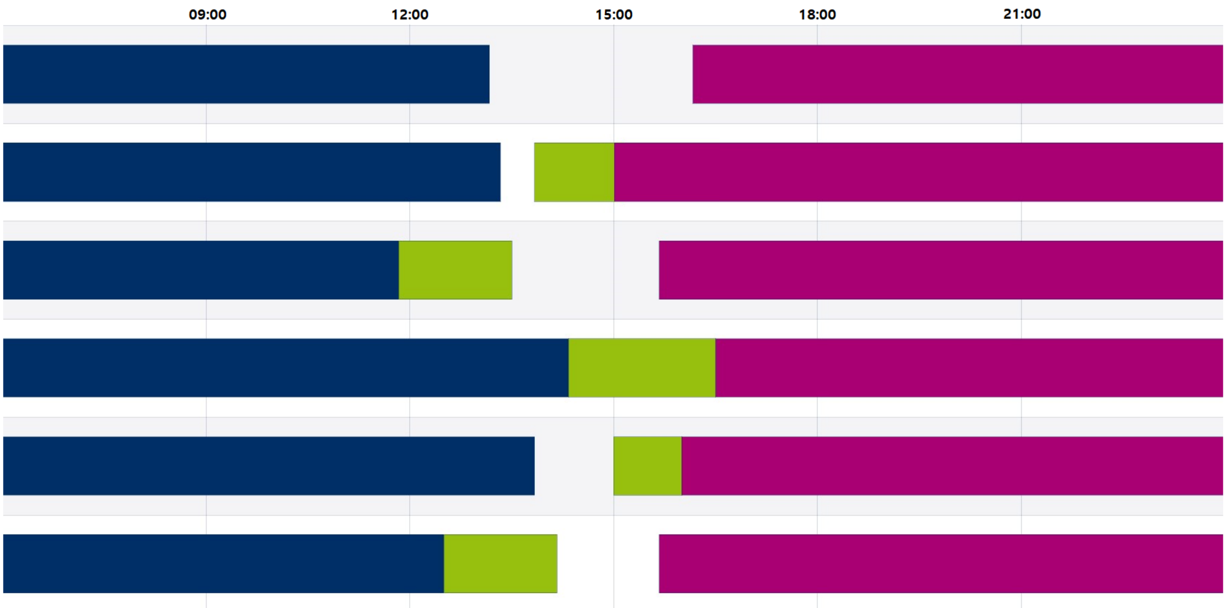

Figure 9.2: Charging Schedule of time between the third and the second set of routes.

When our own model is used, we found the exact same results: two $22kW$ chargers were sufficient for both mid-day and overnight charging. Providing more details about the charging outcomes, or additional charging schedules, will therefore not add to the conclusions that can be drawn.

Part IV

Closure

10
Conclusion

We start this concluding chapter by first summarizing the conclusions draw in each of the four cases that are studied in chapter 9. Then, we can judge the performance and usefulness of the designed models.

10.1. Case I

The goal of the experiments in case I was to gauge how well our routing model performs, when compared to ORTEC's state-of-the-art routing model that would otherwise solve these routing instances. This case is only of secondary importance, as the charging model is the main interest of this project. To make sure that the routing solutions that were used to create charging problems with are of sufficient quality, these comparisons were necessary. The results that we see for the two different fleet types are very different. When using the mixed fleet, we found that the two models differ greatly. Our routing model chose to almost exclusively schedule EVs, while the OHD solutions used a more evenly split fleet. This is fortunate for our research, as it makes the routing problem more challenging for these smaller instances. But, we should be aware that the OHD solver in its current state is more prone to scheduling vehicles without practical range restrictions, than our exact model built with electrical vehicles in mind. When using the fully electric fleet, the results are more comparable. Our routing model managed to find routes that are on average slightly shorter, proving its quality. Therefore, we conclude that our routing model is fit to be used in our combined routing-charging schemes.

10.2. Case II

In case II we instead focused on the charging model. We tested three different sets of instances: the full dataset, a dataset of 100 customers driven by that same fleet, and the same 100 customer dataset driven by a fully electric fleet. Each pair of routes was considered independently, and charging schedules were made under both the assumption that the electricity costs should be as low as possible, and under the assumption that the vehicles should be charged as much as possible. We analyzed the differences between these instances and assumptions, and compared how the different policies described in section 8.2.4 affected the routing solutions found by the OHD solver.

When allowing the vehicles to use their complete range during the shifts (policy 3) instead of only half (policy 1), we found an average cost reduction of 4.6% and a 9.6% decrease of the number of kilometers driven by EVs for the full datasets. For the reduced datasets, this difference is more severe, finding an average cost reduction of 60% and a reduction of 64.6% of non-electrically driven kilometers. Both instances were allowed to schedule the same set of vehicles, so the large difference between the two scenarios is explained mostly by the abundance of electrical vehicles to schedule for the reduced dataset. Either way, this is a significant improvement in both cases. For the fully electric fleet, we instead compared the results of policy (3) with policy (2), that is very similar to policy (3) except that when vehicles are unable to drive a shift, another vehicles is scheduled. Here, we found that we needed a fleet that is 55.5% larger than for policy (3) to complete the shifts. Making the same comparison for the mixed fleet of vehicles, we found that by allowing the vehicles to be charged between shifts, the number of needed vehicles can shrink by 33% up to 50%. Policy (3) assumes that the vehicles can charge during the day, so in order to realize these improvements it is necessary to find feasible charging schedules.

Allowing each of the three instances types access to a realistic set of chargers, we saw that finding such a charging schedule was always possible. For all of these instances, it was possible to charge the complete set of 5 to 54 vehicles enough for their next shift with four chargers (both $22kW$ and $50kW$). To fully charge the vehicles between shifts both during the day and overnight, more chargers were needed. At most, almost as many chargers as vehicles were used, but less than half the number of chargers as there were vehicles was also enough to fully charge the vehicles overnight. These sets of chargers were also able to charge the vehicles in most of the instances in the reduced scenarios fully during the day. In case of the full-sized instances, over 70% of the energy used in the previous shift was recharged before the start of the next shift later that day. Note that this does not mean that the numbers of chargers used in these experiments are the minimum possible amount. We only conclude that a large set of high-capacity chargers is not necessary to reap the benefits of charging the vehicles during the day between shifts.

10.3. Case III

The two models were finally combined in case III. The analysis we performed for this case is similar to that of case II, as the cases are similar. What differs are the engines used to run the the routing instances, the size of those instances and some assumptions about the relationship between the different instances of the same type. In case II, we assumed that all charging instances are independent. In this case, we use model variation 2 described in section 7.3.2, running the routing and charging instances back-to-back. This means that we assume that the vehicles driven in all later instances are the same as the ones driven in the first instance, with the exception of vehicles that are added later. We also tested model variation 3 for the mixed fleet, when an insufficient number of chargers was provided.

Moving from policy (1) to policy (3) also proved to be greatly beneficial for the instances of 50 customers calculated by our routing model. We saw that the costs decreased on average by 71%, finding an average decrease of 11.2% of the total number of vehicles driven. We additionally saw that the number of kilometers driven by non-electric vehicles decreased by an average of 79.9%, meaning that most kilometers are now driven by EVs. When using the fully electric fleet, moving from policy (2) to policy (3), we found that almost no vehicles were able to drive a second shift without charging beforehand. This means that for some individual instances we need to have access to twice as many vehicles to complete the shifts if no in-between charging is allowed. Overall, due to the varying amounts of vehicles needed to complete each set of shifts, the size of the fleet necessary to complete all the tasks decreased by 40% when allowing mid-day charging.

While studying the charging schedules, we considered both the version with and without additional post-processing. While post-processing will always be applied in realistic situations, if necessary, this comparison was made mostly to show what happens when the charging windows are much tighter. We found that with post-processing, all vehicles could be fully charged with four chargers or less for most of the datasets. When the ending or starting times were not considered before creating the routing schedule, the outcomes were worse but not impossibly so. For both the mixed as well as the fully electric fleet, at some point an extra vehicle was needed because the routes were planned so tightly, that even without any charging two shifts would be forced to overlap. For the end-result this did not matter however, as we needed to use extra vehicles regardless, due to later instances that consisted of more shifts. For these instances, we also found that the vehicles could not be charged fully before the start of second shift of the day. It was however possible to recharge over 70% of the energy used during their shifts, which is still the majority. This shows that even when shifts are planned close together, a lot can still be achieved by charging the vehicles during the moments in between.

10.4. Case IV

In case IV we tried out the model variation designed to help determine the right number of vehicles and chargers to purchase. When designing these features, the idea was that by allowing the model to add extra vehicles and chargers at a cost, the solutions found might be more considerate of the costs attached to these purchases. These predictions did not come true, and instead we found that the added complexity of this variation of the routing model only resulted in worse solutions compared to the standard version of the model, when given the same amount of time. The charging model did not have this same issue, but its use can also be questioned in this application. The best set of chargers to install is not just the cheapest possible option of charging the vehicles in theory, but should also factor in human error and other unforeseen circumstances. This is something our model does not account for, making its use limited.

One thing that we did learn from applying this model is that only a small number of slower chargers (about one third of the number of vehicles used) are needed in order to find a feasible charging schedule for the instances that we tested.

10.5. Overview

Overall, the main conclusion we can draw is that even when only a modest set of chargers is available, there is still a lot to achieve by allowing vehicles to be charged between different shifts on the same day. The exact financial benefits completely depend on each scenario, but in our case we found an average reduction of 4.6% when there were no additional EVs to schedule, and up to 71% when this opportunity did exist. More interesting is the reduction of kilometers driven by non-electric vehicles. This decreased on average 9.6% for the full datasets, when there were no additional EVs to schedule, and up to 79% for the reduced datasets, when this was possible.

11 Discussion

11.1. Summary

We start the discussion by providing a brief summary of our key results. In case I we found that our routing model matches up well with ORTEC's solver when only EVs were available, but when vehicles without range restriction were also possible to add, our routing model did add a much larger proportion of EVs. While analyzing the differences between the instances where vehicles were allowed to charge mid-day, and instances where that was not possible, we found that the costs decreased by 4.6% up to 60%, and the share of kilometers driven by non-EVs decreased on average 20% up to 64.6% depending on the instance. Case III showed similar results, resulting in a cost reduction of 71% and a reduction of non-electrically driven kilometers of 79%. We also proved that even when the charging windows were a lot smaller than necessary, all shifts could still be completed with access to a reasonable set of chargers. Finally, case IV taught us that the absolute minimal number of chargers needed for a feasible charging schedule is very low: two $22kW$ chargers were sufficient for up to seven EVs.

11.2. Interpretations and Implications

In this research, we showed that vehicles can be equipped much more efficiently if they are allowed to be charged in-between their shifts during the day, instead of only overnight. As a consequence, the cost as well as the carbon footprint are reduced significantly. The difference this makes is remarkable, especially when considering that the vehicles only need to be charged a little in order to be able to drive longer distances for both shifts. When the charging windows are large enough, which is the case for the datasets that have been used here, we even found that the vehicles can be fully charged after each shift, even when there were twice as many vehicles as chargers. More importantly, the created model allows us to quickly find such solutions, either having minimized the charging cost, or maximized the battery level before departure.

We hope that this research encourages companies to further evaluate the part that electric vehicles play in their business operations. Especially companies that already make use of a (partially) electric fleet stand to benefit greatly from minor improvements in their logistical processes. If it is possible for the vehicles to be stationary for up to a few hours in-between shifts, or even less than that with enough charger availability, our charging model creates schedules that find the most efficient ways in which vehicles can be assigned to both chargers and shifts.

11.3. Limitations and Considerations

This section will describe several limitations of the models we designed and other possible aspects that need to be considered before applying them directly. The biggest limitation of the combined charging and routing model is that it is a combination of two separate optimization models. Even if the separate model pieces returned optimal solutions, we are not guaranteed to have found a global optimum. We initially planned on iteratively applying our models, aiming to find improved solutions. We did not end up (fully) implementing this idea, but do share some of our thoughts in the next section. The alternative would be combining the two models into a single model, solving the entire problem in one go. Theoretically, this could be a good idea and formulating such a problem should be possible, but trying to solve such a problem exactly is not a simple task. In fact, problems based around this idea, such as the Multi-Period

EVRP [134] or the Generalized Periodic EVRP [135], have been studied already. There hasn't been a lot of research done on these variations however, and there are a few apparent downsides. Firstly, the possibilities of including restrictions or adding freedom to the charging problem are quite limited. The level of detail that our charging model can incorporate into the schedules is not seen in these variations, which makes it easier to directly apply these schedules. We also see that these periodic EVRP variations are solved exclusively heuristically. The quality of solutions found heuristically can be very high, but especially due to the complexity of this problem, an optimal solution is not guaranteed. The same of course applies to our dual-model algorithm, but as the charging model is solved exactly, and not complex enough to need to be terminated prematurely, that part of the problem can be solved to optimality. There also already exist very high performing solvers for routing instances. Combining these with this charging model might not yield completely optimal solutions, but based on our experiments we have no reason to believe that directly combining the two models will significantly boost the quality of our solutions. In fact, we would argue that having the charging model separate from the routing model can be more efficient and practical, because then a heuristic does not need to be designed and implemented from scratch. There is definitely value in solving these problems simultaneously, but for our application that is not the most sensible approach.

A smaller limitation of this model can be seen in case IV in section 9.4 of chapter 9. We concluded that when our goal is to determine which chargers should be purchased, our model in its current form is not a suitable method. The reason for this is that it does not consider factors like allowing a charger to (dis)connect to or from vehicles in the middle of the night, which are relevant when finding a schedule that can realistically be applied. Another reason is that there are other factors that come into play when a company determines to install new charging hardware other than the price of this hardware. Installing several fast-chargers could be very desirable, granting the company freedom to quickly charge vehicles in unforeseen circumstances. Our model would only have made this choice if scheduling all tasks was impossible without such a charger due to its high price. We conclude that the charging model is better suited for its original task: finding charging schedules.

We similarly found that the performance of model variation 3, tested in case III in section 9.3, does not to live up to our expectations in the instances tested. While we still believe that this model variation can be very valuable, its use is limited in situations that are too, or insufficiently, constrained. If the set of available vehicles or chargers is too large, the solutions found for this variation do not differ from the solutions found by the standard model variation. When the available set of vehicles or chargers is too strict, routes will instead be driven by ICEVs, which was not the solution we hoped to find using this more sophisticated model variation. Therefore, this model variation only has limited applicability.

Another way in which we can improve our model is to force the model to pick the cheapest possible vehicle when the costs are not immediately of concern. In case IV in section 9.4 of chapter 9 we had no issues, and found that the model only chose to uplift a more expensive vehicle if the route it drove could not have been completed by a cheaper vehicle, but the same thing could not be said about case III in section 9.3. Here, due to the fact that the model did not need to pay for introducing a new vehicle to the solution, sometimes shorter routes were assigned to more expensive vehicles. When giving the results, we instead assigned those routes to a cheaper vehicle, resulting in a cheaper solution of exactly the same quality. In practice, this is no problem, but it should be considered before adopting the schedules that follow from this routing model. A consequence is however that we sometimes find small disparities between energy costs in the charging problem, as the cheaper and more expensive vehicle do have a slightly different energy consumption rate. But, seeing how this energy consumption rate is an estimate to begin with, realistically the energy costs will differ regardless.

We would also like to comment on another factor that affects to what extent the exact output of this model applies to the real situation. As mention earlier this section, the routing and charging model make use of different time systems. The routing model will find continuous values for starting moments, rounded down to minutes, possibly adapted by post-processing, while the charging model makes use of 10 minute long timesteps. The length of these timesteps can of course be changed, but doing so will significantly impact the complexity of the solutions. We currently assume worst-case rounding when the starting and finishing times of shifts and vehicles are exported to the charging model, to make sure that all our schedules are feasible. The downside of this is that in the worst case, the charging window of a vehicle could be short by nearly 20 minutes compared to what it would be in reality, excluding better possible solutions.

A limitation that does significantly impact the quality of the routing solution is the fact that it is an exact model, and therefore cannot solve large instances. We solved this by tasking the model to only solve instances of 50 customers and allowing a runtime of 90 minutes. Despite this, optimal solutions were never found. We did manage to improve the performance of this model to some extent, as described in section 7.2.3, but this still had only limited effect. Other attempts to speed up the model proved fruitless. One thing we attempted was to quickly find a good initial solution, but as that is not the most time-consuming part, we halted those efforts. We also tried to add cuts to the formulation, but this did not end up improving the run-time, so we also did not proceed with that approach.

11.4. Recommendations

In this section we want to provide some inspiration for further research. One thing that can be noted from reading chapter 7 and chapter 8 is that not all modeling functionalities were tested. In fact, some of the more interesting features were not applied during the tests. This was the case for the multi-trip version of the routing model. Implementing this variation was necessary, as the routes calculated by OHD might contain vehicles taking multiple trips, but we did find that for the cases tested by our routing algorithm this was not needed. The battery capacity proved to be a larger bottleneck than the loading capacity of the vehicles. For the charging model we did also not test all implemented features, the most noteworthy of which being non-linear charging. The first reason for this is the additional complexity, that would have made performing the experiments more costly. The second reason is that the non-linear charging functionality depends on several parameters that describe the non-linear charging process of AC-chargers piecewise-linearly. The set of chargers used in our experiments however also contains DC chargers. These chargers have a different charging curve, seen in figure 4.3. This illustrates that for DC-charging different approximations are necessary. There were also a few other smaller features that did not end up in the final experiments, such as variable electricity pricing or vehicles leaving from multiple depots. It would be very interesting to see how these different model variations behave and how they compare to their simpler counterparts.

Section 7.3 also explained and illustrated several different model variations that combined the charging and the routing model. In three of the four described variations, the routing model and the charging model only had limited interaction: the charging model got its input from previous routing and possibly charging outcomes, but the routing models each ran independently. The third model variation does allow information about the charging opportunities to impact the outcome of the routing model, but this is still rather limited. Only routing instances that occur after a period of charging will be affected by this model, meaning that information about later charging or routing opportunities cannot be used. At the beginning of the project, we had several ideas about ways in which the charging and the routing model could interact more, in order to find improved solutions. The main idea was to iteratively calculate the routing and charging model multiple times, allowing information about for example the charging windows and battery requirements of previous and future instances to affect the current instance. In figure 11.1, we illustrated what such a scheme could look like for a single charging instance. This would of course get much more complicated when more charging instances are calculated back-to-back.

The benefit of such a scheme could be that for difficult instances, routing choices in the early, as well as the later, sets of shifts could be re-evaluated depending on the availability of the chargers and charging needs of the vehicles. There are several different scenarios for which such options can be beneficial. Take for example a scenario in which the vehicles driving the first shift of the day all return very late, resulting in very short charging intervals. This might make it impossible for most of those vehicles to drive their next shifts, requiring other vehicles, possibly non-electrical ones, to complete the shifts. In hindsight, a better routing solution would be to instead create slightly more, shorter shifts. These shifts allow the used electric vehicles to return sooner, which means that they can still drive the second set of shifts. Not only does this mean that less shifts will have to be driven by conventional 'back-up' vehicles, we also see that having this ability to improve unfortunate shifts means that less vehicles are needed overall to be able to meet the delivery demands.

We feel that these types of situations would be very interesting to study, and would have liked to do so ourselves within this project. There was however limited applicability for such expansions in the cases that we were studying. Since we were working with real data from a company, the idea was to assume realistic instances overall. This means for example that we do have enough vehicles and chargers available to always be able to visit the scheduled customers. A company is not going to be happy with solutions in which not all customers are visited consistently, even if this means that less charging infrastructure is

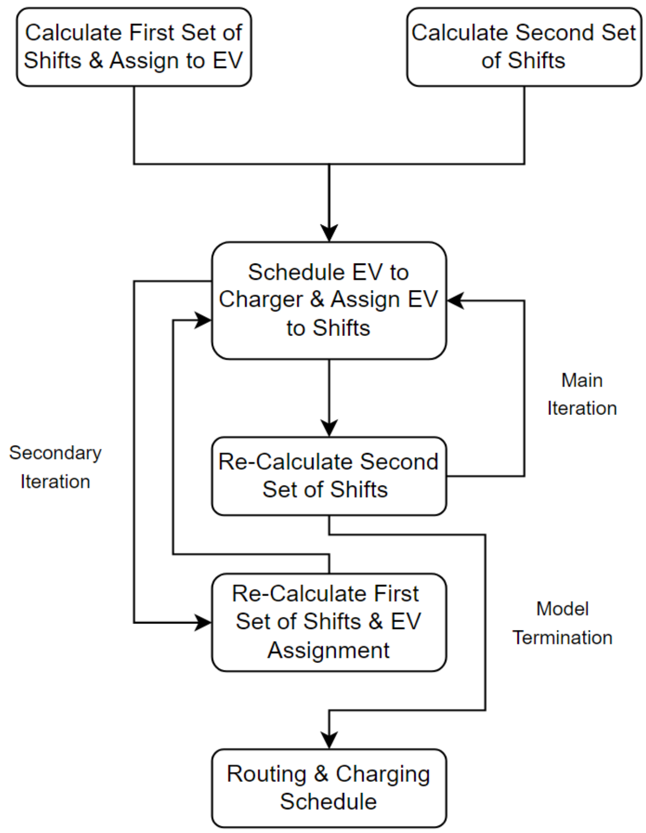

Figure 11.1: Flowchart of a possible model expansion.

necessary. That does not mean that such an expansion has no purpose. This expansion can be very valuable when more challenging instances are encountered regularly. Both for this reason, as well as the limited time available for this project, we decided to not develop this model variation.

References

[1] United Nations, *Paris agreement*, UNTC XXVII 7.d, Dec. 12, 2015.

[2] United Nations, ”What is the Paris Agreement?”, unfccc.int. https://unfccc.int/process-and-meetings/the-paris-agreement (accessed 14 Mar. 2024).

[3] H. Ritchie, P. Rosado, and M. Roser, “Emissions by sector: Where do greenhouse gases come from?” *Our World in Data*, 2020. [Online]. Available: https://ourworldindata.org/emissions-by-sector.

[4] C. Lellis, ””These 21 Companies Are Switching to Electric Vehicle Fleets”, www.perillon.com. https://www.perillon.com/blog/21-companies-switching-to-electric-vehicle-fleets (accessed 15 Mar. 2024)”.

[5] İ. Küçükoğlu, R. Dewil, and D. Cattrysse, “The electric vehicle routing problem and its variations: A literature review,” *Computers & Industrial Engineering*, vol. 161, p. 107 650, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0360835221005544.

[6] EVBox, ”What do workplace EV chargers cost?”, blog.evbox.com. https://blog.evbox.com/uk-en/ev-workplace-charger-costs (accessed 15 Mar. 2024).

[7] R. Kannan and C. L. Monma, “On the computational complexity of integer programming problems,” in *Optimization and Operations Research*, Berlin, Heidelberg: Springer Berlin Heidelberg, 1978, pp. 161–172.

[8] M. Velednitsky, “Short combinatorial proof that the DFJ polytope is contained in the MTZ polytope for the asymmetric traveling salesman problem,” *CoRR*, vol. abs/1805.06997, 2018. arXiv: 1805.06997. [Online]. Available: http://arxiv.org/abs/1805.06997.

[9] M. Grötschel and G. L. Nemhauser, “George dantzig’s contributions to integer programming,” *Discrete Optimization*, vol. 5, no. 2, pp. 168–173, 2008, In Memory of George B. Dantzig. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1572528607000412.

[10] L. Kerckhoffs, *Models, solutions and relaxations of the asymmetrical capacitated vehicle routing problem*, BSc. Thesis, EEMCS, Delft University of Technology, Delft, 2007. Available: https://repository.tudelft.nl/islandora/object/uuid:03af3d1b-98d8-4c14-99ff-a448b4f4b2d0?collection =education.

[11] L. A. Wolsey, *Integer Programming*. New York: Wiley, 1998.

[12] Y. Lim and Y. J. Yu, “A width-recursive depth-first tree search approach for the design of discrete coefficient perfect reconstruction lattice filter bank,” *Circuits and Systems II: Analog and Digital Signal Processing, IEEE Transactions on*, vol. 50, pp. 257–266, Jul. 2003.

[13] D. R. Morrison, S. H. Jacobson, J. J. Sauppe, and E. C. Sewell, “Branch-and-bound algorithms: A survey of recent advances in searching, branching, and pruning,” *Discrete Optimization*, vol. 19, pp. 79–102, 2016. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1572528616000062.

[14] T. Achterberg, T. Koch, and A. Martin, “Branching rules revisited,” *Operations Research Letters*, vol. 33, no. 1, pp. 42–54, 2005. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0167637704000501.

[15] G. Clarke and J. Wright, “Scheduling of vehicles from a central depot to a number of delivery points,” *Operations Research*, vol. 12, no. 4, pp. 519–643, 1964.

[16] G. B. Dantzig and J. H. Ramser, “The truck dispatching problem,” *Management science*, vol. 6, no. 1, pp. 80–91, 1959.

[17] K. Braekers, K. Ramaekers, and I. Van Nieuwenhuyse, “The vehicle routing problem: State of the art classification and review,” *Computers & Industrial Engineering*, vol. 99, pp. 300–313, 2016. [Online]. Available: `https://www.sciencedirect.com/science/article/pii/S0360835215004775`.

[18] Y. Marinakis, M. Marinaki, and A. Migdalas, “Variants and formulations of the vehicle routing problem,” in *Open Problems in Optimization and Data Analysis*, P. M. Pardalos and A. Migdalas, Eds. Cham: Springer International Publishing, 2018, pp. 91–127. [Online]. Available: `https://doi.org/10.1007/978-3-319-99142-9_7`.

[19] Wikipedia contributors, ”Dynamic programming”, wikipedia.com. https://en.wikipedia.org/w/index.php ?title=Dynamic_programming&oldid=1172005607 (accessed 11 Mar. 2024).

[20] N. Christofides, A. Mignozzi, and P. Toth, “The vehicle routing problem,” *Combinatorial optimization*, pp. 315–338, 1979.

[21] P. Munari, T. Dollevoet, and R. Spliet, “A generalized formulation for vehicle routing problems,” Sep. 2017.

[22] C. E. Miller, A. W. Tucker, and R. A. Zemlin, “Integer programming formulation of traveling salesman problems,” *J. ACM*, vol. 7, no. 4, pp. 326–329, 1960. [Online]. Available: `https://doi.org/10.1145/321043.321046`.

[23] P. Toth and D. Vigo, “The vehicle routing problem,” 2002. [Online]. Available: `https://epubs.siam.org/doi/abs/10.1137/1.9780898718515`.

[24] M. Desrochers and G. Laporte, “Improvements and extensions to the miller-tucker-zemlin subtour elimination constraints,” *Operations Research Letters*, vol. 10, no. 1, pp. 27–36, 1991. [Online]. Available: `https://www.sciencedirect.com/science/article/pii/0167637791900832`.

[25] S. Irnich, P. Toth, and D. Vigo, “Chapter 1: The family of vehicle routing problems,” in Nov. 2014, pp. 1–33.

[26] W. W. Garvin, H. W. Crandall, J. B. John, and R. A. Spellman, “Applications of linear programming in the oil industry,” *Management Science*, vol. 3, pp. 407–430, 1957. [Online]. Available: `https://api.semanticscholar.org/CorpusID:154332186`.

[27] B. Gavish and S. C. Graves, “The travelling salesman problem and related problems,” 1978. [Online]. Available: `https://api.semanticscholar.org/CorpusID:14353795`.

[28] B. Gavish and S. C. Graves, “Scheduling and routing in transportation and distributions systems: Formulations and new relaxations,” University of Rochester, 1982. [Online]. Available: `https://api.semanticscholar.org/CorpusID:14353795`.

[29] R. Baldacci, E. Hadjiconstantinou, and A. Mingozzi, “An exact algorithm for the capacitated vehicle routing problem based on a two-commodity network flow formulation,” *Operations Research*, vol. 52, pp. 723–738, Oct. 2004.

[30] M. Balinski and R. E. Quandt, “On an integer program for a delivery problem,” *Operations Research*, vol. 12, no. 2, pp. 300–304, 1964. [Online]. Available: `https://EconPapers.repec.org/RePEc:inm:oropre:v:12:y:1964:i:2:p:300-304`.

[31] R. Basso, P. Lindroth, B. Kulcsár, and B. Egardt, “Traffic aware electric vehicle routing,” in *2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC)*, 2016, pp. 416–421.

[32] M. Soysal, M. Çimen, and S. Belbağ, “Pickup and delivery with electric vehicles under stochastic battery depletion,” *Computers & Industrial Engineering*, vol. 146, p. 106 512, 2020. [Online]. Available: `https://www.sciencedirect.com/science/article/pii/S0360835220302461`.

[33] H. Kopfer and B. Vornhusen, “Energy vehicle routing problem for differently sized and powered vehicles,” *Journal of Business Economics*, vol. 89, pp. 793–821, 2019.

[34] J. Chen, M. Qi, and L. Miao, “The electric vehicle routing problem with time windows and battery swapping stations,” in *2016 IEEE International Conference on Industrial Engineering and Engineering Management (IEEM)*, 2016, pp. 712–716.

[35] D. Wang, H. Zhou, and R. Feng, “A two-echelon vehicle routing problem involving electric vehicles with time windows,” in *Journal of Physics: Conference Series*, IOP Publishing, vol. 1324, 2019, p. 012 071.

[36] Y. Xiao, X. Zuo, I. Kaku, S. Zhou, and X. Pan, “Development of energy consumption optimization model for the electric vehicle routing problem with time windows,” *Journal of Cleaner Production*, vol. 225, pp. 647–663, 2019. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S095965261931039X.

[37] J. Lu, Y. Chen, J.-K. Hao, and R. He, “The time-dependent electric vehicle routing problem: Model and solution,” *Expert Systems with Applications*, vol. 161, p. 113 593, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0957417420304176.

[38] S. R. Kancharla and G. Ramadurai, “Electric vehicle routing problem with non-linear charging and load-dependent discharging,” *Expert Systems with Applications*, vol. 160, p. 113 714, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0957417420305388.

[39] A. Froger, J. E. Mendoza, O. Jabali, and G. Laporte, “Improved formulations and algorithmic components for the electric vehicle routing problem with nonlinear charging functions,” *Computers & Operations Research*, vol. 104, pp. 256–294, 2019. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0305054818303253.

[40] C. Lee, “An exact algorithm for the electric-vehicle routing problem with nonlinear charging time,” *Journal of the Operational Research Society*, vol. 72, no. 7, pp. 1461–1485, 2021. [Online]. Available: https://doi.org/10.1080/01605682.2020.1730250.

[41] H. Tahami, G. Rabadi, and M. Haouari, “Exact approaches for routing capacitated electric vehicles,” *Transportation Research Part E: Logistics and Transportation Review*, vol. 144, p. 102 126, 2020.

[42] A. Ceselli, Á. Felipe, M. T. Ortuño, and G. Righini Giovanniand Tirado, “A branch-and-cut-and-price algorithm for the electric vehicle routing problem with multiple technologies,” *Operations Research Forum*, vol. 2, no. 1, p. 8, Jan. 2021. [Online]. Available: https://doi.org/10.1007/s43069-020-00052-x.

[43] P. Munari, A. Moreno, J. De La Vega, D. Alem, J. Gondzio, and R. Morabito, “The robust vehicle routing problem with time windows: Compact formulation and branch-price-and-cut method,” *Transportation Science*, vol. 53, no. 4, pp. 1043–1066, 2019. [Online]. Available: https://doi.org/10.1287/trsc.2018.0886.

[44] L. Costa, C. Contardo, and G. Desaulniers, “Exact branch-price-and-cut algorithms for vehicle routing,” *Transportation Science*, vol. 53, no. 4, pp. 946–985, 2019. [Online]. Available: https://doi.org/10.1287/trsc.2018.0878.

[45] M. Barth, T. Younglove, and G. Scora, “Development of a heavy-duty diesel modal emissions and fuel consumption model,” *Institute of Transportation Studies, UC Berkeley, Institute of Transportation Studies, Research Reports, Working Papers, Proceedings*, Jan. 2005.

[46] D. Goeke and M. Schneider, “Routing a mixed fleet of electric and conventional vehicles,” *European Journal of Operational Research*, vol. 245, no. 1, pp. 81–99, 2015. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0377221715000697.

[47] J. Asamer, A. Graser, B. Heilmann, and M. Ruthmair, “Sensitivity analysis for energy demand estimation of electric vehicles,” *Transportation Research Part D: Transport and Environment*, vol. 46, pp. 182–199, 2016. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1361920915300250.

[48] G. Sandrini, D. Chindamo, and M. Gadola, “Regenerative braking logic that maximizes energy recovery ensuring the vehicle stability,” *Energies*, vol. 15, no. 16, 2022. [Online]. Available: https://www.mdpi.com/1996-1073/15/16/5846.

[49] M. Ehsani, Y. Gao, and A. Emadi, *Modern Electric, Hybrid Electric, and Fuel Cell Vehicles: Fundamentals, Theory, and Design (second ed.)* CRC press, 2009, ch. 13.

[50] S. Pelletier, O. Jabali, and G. Laporte, "The electric vehicle routing problem with energy consumption uncertainty," *Transportation Research Part B: Methodological*, vol. 126, pp. 225–255, 2019. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0191261519300153.

[51] R. Basso, B. Kulcsár, B. Egardt, P. Lindroth, and I. Sanchez-Diaz, "Energy consumption estimation integrated into the electric vehicle routing problem," *Transportation Research Part D: Transport and Environment*, vol. 69, pp. 141–167, 2019. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1361920918304760.

[52] S. Hulagu and H. B. Celikoglu, "Electric vehicle location routing problem with vehicle motion dynamics-based energy consumption and recovery," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 8, pp. 10 275–10 286, 2022.

[53] M. Keskin and B. Çatay, "Partial recharge strategies for the electric vehicle routing problem with time windows," *Transportation Research Part C: Emerging Technologies*, vol. 65, pp. 111–127, 2016. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0968090X16000322.

[54] R. Roberti and M. Wen, "The electric traveling salesman problem with time windows," *Transportation Research Part E: Logistics and Transportation Review*, vol. 89, pp. 32–52, 2016. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1366554516000181.

[55] S. Yang, L. Ning, L. C. Tong, and P. Shang, "Optimizing electric vehicle routing problems with mixed backhauls and recharging strategies in multi-dimensional representation network," *Expert Systems with Applications*, vol. 176, p. 114 804, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0957417421002451.

[56] G. Hiermann, J. Puchinger, S. Ropke, and R. F. Hartl, "The electric fleet size and mix vehicle routing problem with time windows and recharging stations," *European Journal of Operational Research*, vol. 252, no. 3, pp. 995–1018, 2016. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0377221716000837.

[57] M. Zhao and Y. Lu, "A heuristic approach for a real-world electric vehicle routing problem," *Algorithms*, vol. 12, no. 2, 2019. [Online]. Available: https://www.mdpi.com/1999-4893/12/2/45.

[58] J. Barco Jiménez, A. Guerra, L. Muñoz, and N. Quijano, "Optimal routing and scheduling of charge for electric vehicles: A case study," *Mathematical Problems in Engineering*, vol. 2017, Nov. 2017.

[59] L. Li, T. Li, K. Wang, S. Gao, Z. Chen, and L. Wang, "Heterogeneous fleet electric vehicle routing optimization for logistic distribution with time windows and simultaneous pick-up and delivery service," in *2019 16th International Conference on Service Systems and Service Management (ICSSSM)*, 2019, pp. 1–6.

[60] J. P. Futalef, D. Munoz-Carpintero, H. Rozas, and M. Orchard, "An evolutionary algorithm for the electric vehicle routing problem with battery degradation and capacitated charging stations," in *Annual Conference of the PHM Society*, vol. 12, 2020, pp. 9–9.

[61] Y. J. Kim and B. D. Chung, "Energy consumption optimization for the electric vehicle routing problem with state-of-charge-dependent discharging rates," *Journal of Cleaner Production*, vol. 385, p. 135 703, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0959652622052775.

[62] R. Basso, B. Kulcsár, and I. Sanchez-Diaz, "Electric vehicle routing problem with machine learning for energy prediction," *Transportation Research Part B: Methodological*, vol. 145, pp. 24–55, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0191261520304549.

[63] S. Zhang, M. Chen, W. Zhang, and X. Zhuang, "Fuzzy optimization model for electric vehicle routing problem with time windows and recharging stations," *Expert Systems with Applications*, vol. 145, p. 113 123, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0957417419308401.

[64] Citroën, "Citroën ë-Jumpy technische gegevens", business.citroen.nl. https://business.citroen.nl/content/dam/citroen/netherlands/b2b/local-content/brochures-prijslijsten/model/e-jumpy/wk0421_technische-specificaties-e-jumpy.369216.pdf (accessed 11 Mar. 2024).

[65] The Engineering Toolbox, ”Rolling Resistance”, engineeringtoolbox.com. https://www.engineeringtoolbox.com/rolling-friction-resistance-d_1303.html (accessed 11 Mar. 2024).

[66] H. Chowdhury, R. Juwono, M. Zaid, R. Islam, B. Loganathan, and F. Alam, “An experimental study on of the effect of various deflectors used for light trucks in indian subcontinent,” *Energy Procedia*, vol. 160, pp. 34–39, Feb. 2019.

[67] C. Fiori, K. Ahn, and H. A. Rakha, “Power-based electric vehicle energy consumption model: Model development and validation,” *Applied Energy*, vol. 168, pp. 257–268, 2016. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S030626191630085X.

[68] X. Yuan, C. Zhang, G. Hong, X. Huang, and L. Li, “Method for evaluating the real-world driving energy consumptions of electric vehicles,” *Energy*, vol. 141, pp. 1955–1968, 2017. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0360544217319928.

[69] X. Wu, D. Freese, A. Cabrera, and W. A. Kitch, “Electric vehicles’ energy consumption measurement and estimation,” *Transportation Research Part D: Transport and Environment*, vol. 34, pp. 52–67, 2015. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1361920914001485.

[70] O. Travesset-Baro, M. Rosas-Casals, and E. Jover, “Transport energy consumption in mountainous roads. a comparative case study for internal combustion engines and electric vehicles in andorra,” *Transportation Research Part D: Transport and Environment*, vol. 34, pp. 16–26, 2015. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1361920914001424.

[71] T. Yuksel and J. J. Michalek, “Effects of regional temperature on electric vehicle efficiency, range, and emissions in the united states,” *Environmental Science & Technology*, vol. 49, no. 6, pp. 3974–3980, Mar. 2015.

[72] G. M. Fetene, S. Kaplan, S. L. Mabit, A. F. Jensen, and C. G. Prato, “Harnessing big data for estimating the energy consumption and driving range of electric vehicles,” *Transportation Research Part D: Transport and Environment*, vol. 54, pp. 1–11, 2017. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1361920917303188.

[73] Y. Xie, Y. Li, Z. Zhao, *et al.*, “Microsimulation of electric vehicle energy consumption and driving range,” *Applied Energy*, vol. 267, p. 115 081, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0306261920305936.

[74] P. Iora and L. Tribioli, “Effect of ambient temperature on electric vehicles’ energy consumption and range: Model definition and sensitivity analysis based on nissan leaf data,” *World Electric Vehicle Journal*, vol. 10, no. 1, 2019. [Online]. Available: https://www.mdpi.com/2032-6653/10/1/2.

[75] A. Donkers, D. Yang, and M. Viktorović, “Influence of driving style, infrastructure, weather and traffic on electric vehicle performance,” *Transportation Research Part D: Transport and Environment*, vol. 88, p. 102 569, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1361920920307562.

[76] K. Liu, T. Yamamoto, and T. Morikawa, “Impact of road gradient on energy consumption of electric vehicles,” *Transportation Research Part D: Transport and Environment*, vol. 54, pp. 74–81, 2017. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1361920917303887.

[77] K. N. Genikomsakis and G. Mitrentsis, “A computationally efficient simulation model for estimating energy consumption of electric vehicles in the context of route planning applications,” *Transportation Research Part D: Transport and Environment*, vol. 50, pp. 98–118, 2017. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1361920915302881.

[78] Y. Al-Wreikat, C. Serrano, and J. R. Sodré, “Driving behaviour and trip condition effects on the energy consumption of an electric vehicle under real-world driving,” *Applied Energy*, vol. 297, p. 117 096, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0306261921005444.

[79] M. De Gennaro, E. Paffumi, G. Martini, *et al.*, “Experimental investigation of the energy efficiency of an electric vehicle in different driving conditions,” in *SAE 2014 World Congress & Exhibition*, SAE International, Apr. 2014.

[80] S. Mamarikas, S. Doulgeris, Z. Samaras, and L. Ntziachristos, “Traffic impacts on energy consumption of electric and conventional vehicles,” *Transportation Research Part D: Transport and Environment*, vol. 105, p. 103 231, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S136192092200061X.

[81] G. Duarte, R. Varella, G. Gonçalves, and T. Farias, “Effect of battery state of charge on fuel use and pollutant emissions of a full hybrid electric light duty vehicle,” *Journal of Power Sources*, vol. 246, pp. 377–386, 2014. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0378775313013128.

[82] S. El Amrani, M. Chennani, and D. Belkhayat, “Comparative study of electric vehicle energy consumption between trunk roads and highways,” in *2019 7th International Renewable and Sustainable Energy Conference (IRSEC)*, 2019, pp. 1–7.

[83] T. Jonas, T. Wilde, C. Hunter, and G. Macht, “The impact of road types on the energy consumption of electric vehicles,” *Journal of Advanced Transportation*, vol. 2022, Jun. 2022.

[84] A. Doyle and T. Muneer, “Energy consumption and modelling of the climate control system in the electric vehicle,” *Energy Exploration & Exploitation*, vol. 37, pp. 519–543, Oct. 2018.

[85] M. Mruzek, I. Gajdáč, L. Kucera, and D. Barta, “Analysis of parameters influencing electric vehicle range,” *Procedia Engineering*, vol. 134, pp. 165–174, Dec. 2016.

[86] Electreon, ”Projects”, electreon.com. https://electreon.com/projects (accessed 11 Mar. 2024).

[87] J. Ward, ”Five-minute recharge: China’s Nio rolling out new electric-car battery swap stations in Europe and UK”, drive.com.au. https://www.drive.com.au/news/new-nio-battery-swap-stations-europe (accessed 11 Mar. 2024).

[88] R. Raeesi and K. G. Zografos, “The electric vehicle routing problem with time windows and synchronised mobile battery swapping,” *Transportation Research Part B: Methodological*, vol. 140, pp. 101–129, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0191261520303593.

[89] J. Li, F. Wang, and Y. He, “Electric vehicle routing problem with battery swapping considering energy consumption and carbon emissions,” *Sustainability*, vol. 12, no. 24, 2020. [Online]. Available: https://www.mdpi.com/2071-1050/12/24/10537.

[90] M. Meng and Y. Ma, “Route optimization of electric vehicle considering soft time windows and two ways of power replenishment,” *Advances in Operations Research*, vol. 2020, May 2020. [Online]. Available: https://doi.org/10.1155/2020/5612872.

[91] InductEV, inductev.com. https://www.inductev.com (accessed 11 Mar. 2024).

[92] WAVE Charging, wavecharging.com. https://wavecharging.com (accessed 11 Mar. 2024).

[93] WiTricity, witricity.com. https://witricity.com/ (accessed 11 Mar. 2024).

[94] Plugless Power, pluglesspower.com. https://pluglesspower.com (accessed 11 Mar. 2024).

[95] A. A. Mohamed, A. A. Shaier, H. Metwally, and S. I. Selem, “A comprehensive overview of inductive pad in electric vehicles stationary charging,” *Applied Energy*, vol. 262, p. 114 584, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0306261920300969.

[96] E. Apostolaki-Iosifidou, P. Codani, and W. Kempton, “Measurement of power loss during electric vehicle charging and discharging,” *Energy*, vol. 127, pp. 730–742, 2017. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0360544217303730.

[97] H. Shareef, M. M. Islam, and A. Mohamed, “A review of the stage-of-the-art charging technologies, placement methodologies, and impacts of electric vehicles,” *Renewable and Sustainable Energy*

Reviews, vol. 64, pp. 403–420, 2016. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1364032116302568.

[98] U.S. Department of Transportation, "Charger Types and Speeds", transportation.gov. https://www.transportation.gov/rural/ev/toolkit/ev-basics/charging-speeds (accessed 11 Mar. 2024).

[99] EVBox, "EV charging: the difference between AC and DC", blog.evbox.com. https://blog.evbox.com/difference-between-ac-and-dc (accessed 11 Mar. 2024).

[100] A. Montoya, C. Guéret, J. E. Mendoza, and J. G. Villegas, "The electric vehicle routing problem with nonlinear charging function," *Transportation Research Part B: Methodological*, vol. 103, pp. 87–110, 2017, Green Urban Transportation. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0191261516304556.

[101] S. Pelletier, O. Jabali, G. Laporte, and M. Veneroni, "Battery degradation and behaviour for electric vehicles: Review and numerical analyses of several models," *Transportation Research Part B: Methodological*, vol. 103, pp. 158–187, 2017, Green Urban Transportation. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0191261516303794.

[102] L. Lam, *A Practical Circuit-based Model for State of Health Estimation of Li-ion Battery Cells in Electric Vehicles*, MSc. Thesis, EEMCS, Delft University of Technology, Delft, 2011. Available: https://repository.tudelft.nl/islandora/object/uuid%3Aa7446a0a-4c29-4c68-bbe2-273f68f85ed7.

[103] S. Hemavathi and A. Shinisha, "A study on trends and developments in electric vehicle charging technologies," *Journal of Energy Storage*, vol. 52, p. 105 013, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2352152X22010179.

[104] H. Hõimoja, A. Rufer, G. Dziechciaruk, and A. Vezzini, "An ultrafast ev charging station demonstrator," in *International Symposium on Power Electronics Power Electronics, Electrical Drives, Automation and Motion*, 2012, pp. 1390–1395.

[105] Wikipedia contributors, "Recovery effect", wikipedia.com. https://en.wikipedia.org/w/index.php?title=Recovery_effect&oldid=1076288481 (accessed 11 Mar. 2024).

[106] X. Wu, C. Hu, J. Du, and J. Sun, "Multistage cc-cv charge method for li-ion battery," *Mathematical Problems in Engineering*, vol. 2015, p. 294 793, Oct. 2015. [Online]. Available: https://doi.org/10.1155/2015/294793.

[107] M. Bruglieri, A. Colorni, and A. Luè, "The vehicle relocation problem for the one-way electric vehicle sharing: An application to the milan case," *Procedia - Social and Behavioral Sciences*, vol. 111, pp. 18–27, 2014. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1877042814000354.

[108] M. Uhrig, L. Weiß, M. Suriyah, and T. Leibfried, "E-mobility in car parks–guidelines for charging infrastructure expansion planning and operation based on stochastic simulations," in *EVS28 International Electric Vehicle Symposium and Exhibition*, 2015, pp. 1–12.

[109] K. Erdoğdu and K. Karabulut, "Distance and energy consumption minimization in electric traveling salesman problem with time windows," in *2020 7th International Conference on Electrical and Electronics Engineering (ICEEE)*, 2020, pp. 160–164.

[110] K. S. Abdallah and Y. Adel, "Electric vehicles routing problem with variable speed and time windows," Industry, Engineering & Conference Management Systems Conference, 2020.

[111] H. Li, Z. Li, L. Cao, R. Wang, and M. Ren, "Research on optimization of electric vehicle routing problem with time window," *IEEE Access*, vol. 8, pp. 146 707–146 718, 2020.

[112] K. E. C. Booth and J. C. Beck, "A constraint programming approach to electric vehicle routing with time windows," in *Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, L.-M. Rousseau and K. Stergiou, Eds., Cham: Springer International Publishing, 2019, pp. 129–145.

[113] İ. Küçükoğlu, R. Dewil, and D. Cattrysse, "Hybrid simulated annealing and tabu search method for the electric travelling salesman problem with time windows and mixed charging rates," *Expert Systems*

with Applications, vol. 134, pp. 279–303, 2019. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0957417419303732.

[114] R. G. Conrad and M. A. Figliozzi, “The recharging vehicle routing problem,” in *Proceedings of the 2011 industrial engineering research conference*, IISE Norcross, GA, vol. 8, 2011.

[115] H. Mao, J. Shi, Y. Zhou, and G. Zhang, “The electric vehicle routing problem with time windows and multiple recharging options,” *IEEE Access*, vol. 8, pp. 114 864–114 875, 2020.

[116] S. Chen, C. Yao, H. Miao, Z. Li, and Z. Yang, “Chance constrained routing and charging problem of electric vehicles,” in *2022 China Automation Congress (CAC)*, 2022, pp. 3117–3122.

[117] C. Yao, S. Chen, M. Salazar, and Z. Yang, “Incentive-aware electric vehicle routing problem: A bi-level model and a joint solution algorithm,” in *2022 American Control Conference (ACC)*, 2022, pp. 4662–4667.

[118] S. Karakatič, “Optimizing nonlinear charging times of electric vehicle routing with genetic algorithm,” *Expert Systems with Applications*, vol. 164, p. 114 039, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0957417420308083.

[119] Ç. Koç, O. Jabali, J. E. Mendoza, and G. Laporte, “The electric vehicle routing problem with shared charging stations,” *International Transactions in Operational Research*, vol. 26, no. 4, pp. 1211–1243, 2019. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1111/itor.12620.

[120] J. Lindgren and P. D. Lund, “Effect of extreme temperatures on battery charging and performance of electric vehicles,” *Journal of Power Sources*, vol. 328, pp. 37–45, 2016. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0378775316308941.

[121] Y. Motoaki, W. Yi, and S. Salisbury, “Empirical analysis of electric vehicle fast charging under cold temperatures,” *Energy Policy*, vol. 122, pp. 162–168, 2018. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0301421518304828.

[122] M. Roelofs and J. Bisschop, *AIMMS The User’s Guide*, 2023. [Online]. Available: https://documentation.aimms.com/_downloads/AIMMS_user.pdf.

[123] Gurobi Optimization, LLC, *Gurobi Optimizer Reference Manual*, 2024. [Online]. Available: https://www.gurobi.com.

[124] N. A. Wouda, L. Lan, and W. Kool, “PyVRP: a high-performance VRP solver package,” *INFORMS Journal on Computing*, 2024.

[125] J. Lin, W. Zhou, and O. Wolfson, “Electric vehicle routing problem,” *Transportation Research Procedia*, vol. 12, pp. 508–521, 2016, Tenth International Conference on City Logistics 17-19 June 2015, Tenerife, Spain. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2352146516000089.

[126] A. Olivera and O. Viera, “Adaptive memory programming for the vehicle routing problem with multiple trips,” *Computers & Operations Research*, vol. 34, no. 1, pp. 28–47, 2007. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0305054805001346.

[127] L. C. Coelho, J.-P. Gagliardi, J. Renaud, and A. Ruiz, “Solving the vehicle routing problem with lunch break arising in the furniture delivery industry,” *The Journal of the Operational Research Society*, vol. 67, no. 5, pp. 743–751, 2016. [Online]. Available: http://www.jstor.org/stable/44988090 (visited on 11/22/2023).

[128] European Labour Authority, ”Driving and Resting Time Rules”, ela.europa.eu https://www.ela.europa.eu/sites/default/files/2022-08/Driving_and_Resting_Times_English.pdf (accessed 11 Mar. 2024).

[129] ANWB, ”Actieradius van elektrische auto’s: zo werkt het”, anwb.nl. https://www.anwb.nl/auto/elektrisch-rijden/actieradius-en-verbruik (accessed 11 Mar. 2024).

[130] WLTP, ”What is WLTP and how does it work?”, wltpfacts.eu. https://www.wltpfacts.eu/what-is-wltp-how-will-it-work (accessed 11 Mar. 2024).

[131] Toyota, "Proace ElectricLive Worker", toyota.nl. https://www.toyota.nl/modellen/proace-electric/build?path=engine/4e532c9d-f6b7-4d24-a0e8-bab43c5f18fb/0054f775-24cc-4323-a64b-c1615aefd585/32e71948-ca3e-4f5a-b387-2ec8672b1afa# (accessed 11 Mar. 2024).

[132] World Economics, "Temperature Data", worldeconomics.com. https://www.worldeconomics.com/Indicator-Data/ESG/Environment/Temperatures (accessed 11 Mar. 2024).

[133] EasySwitch.nl, "Stroomprijs per kWh", easyswitch.nl. https://www.easyswitch.nl/energieprijzen/ (accessed Nov. 2023).

[134] L. C. Echeverri, A. Froger, J. E. Mendoza, and E. Neron, "A matheuristic for the Multi-period Electric Vehicle Routing Problem," in *13th Metaheuristics International Conference*, Cartagena de Indias, Colombia, Jul. 2019. [Online]. Available: `https://hal.science/hal-02279106`.

[135] T. Oulad Kouider, W. Ramdane Cherif-Khettaf, and A. Oulamara, "Metaheuristics for the generalised periodic electric vehicle routing problem," in *Computational Logistics*, Cham: Springer International Publishing, 2019, pp. 219–232.

[136] N. Mladenović and P. Hansen, "Variable neighborhood search," *Computers & Operations Research*, vol. 24, no. 11, pp. 1097–1100, 1997. [Online]. Available: `https://www.sciencedirect.com/science/article/pii/S0305054897000312`.

[137] D. Pisinger and S. Røpke, "Large neighborhood search," in *Handbook of Metaheuristics*, M. Gendreau, Ed., 2nd ed. Springer, 2010, pp. 399–420. [Online]. Available: `https://backend.orbit.dtu.dk/ws/portalfiles/portal/5293785/Pisinger.pdf`.

[138] F. Glover, "Tabu search—part i," *ORSA Journal on Computing*, vol. 1, no. 3, pp. 190–206, 1989. [Online]. Available: `https://doi.org/10.1287/ijoc.1.3.190`.

[139] M. Dorigo, "Ant colony optimization," *Scholarpedia*, vol. 2, no. 3, p. 1461, 2007, revision #90969.

[140] B. Eksioglu, A. V. Vural, and A. Reisman, "The vehicle routing problem: A taxonomic review," *Computers & Industrial Engineering*, vol. 57, no. 4, pp. 1472–1483, 2009. [Online]. Available: `https://www.sciencedirect.com/science/article/pii/S0360835209001405`.

[141] D. Goeke, "Granular tabu search for the pickup and delivery problem with time windows and electric vehicles," *European Journal of Operational Research*, vol. 278, no. 3, pp. 821–836, 2019. [Online]. Available: `https://www.sciencedirect.com/science/article/pii/S0377221719304126`.

[142] Y.-H. Jia, Y. Mei, and M. Zhang, "A bilevel ant colony optimization algorithm for capacitated electric vehicle routing problem," *IEEE Transactions on Cybernetics*, vol. 52, no. 10, pp. 10 855–10 868, 2021.

[143] W. Li-ying and S. Yuan-bin, "Multiple charging station location-routing problem with time window of electric vehicle.," *Journal of Engineering Science & Technology Review*, vol. 8, no. 5, 2015.

[144] N. Ding, R. Batta, C. Kwon, *et al.*, "Conflict-free electric vehicle routing problem with capacitated charging stations and partial recharge," *SUNY, Buffalo*, 2015. [Online]. Available: `https://www.acsu.buffalo.edu/~batta/Nan%20Ding.pdf`.

[145] Z. C. Taşkin, "Benders decomposition," in *Wiley Encyclopedia of Operations Research and Management Science*. John Wiley & Sons, Ltd, 2011. [Online]. Available: `https://onlinelibrary.wiley.com/doi/abs/10.1002/9780470400531.eorms0104`.

[146] A. Aich, "A Brief Review of Alternating Direction Method Of Multipliers (ADMM)", abhishekaich27.github.io. https://abhishekaich27. github.io/data/WriteUps/ADMM_Review.pdf (accessed 11 Mar. 2024).

[147] Wikipedia contributors, "Bellman–Ford algorithm", wikipedia.com. https://en.wikipedia.org/w/index.php?title=Bellman%E2%80%93Ford_algorithm&oldid=1173688567 (accessed 11 Mar. 2024).

[148] C. Yao, S. Chen, and Z. Yang, "Joint routing and charging problem of multiple electric vehicles: A fast optimization algorithm," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 7, pp. 8184–8193, 2022.

[149] S. Schoenberg and F. Dressler, “Reducing waiting times at charging stations with adaptive electric vehicle route planning,” *IEEE Transactions on Intelligent Vehicles*, vol. 8, no. 1, pp. 95–107, 2023.

[150] Wikipedia contributors, ”Contraction hierarchies”, wikipedia.com. https://en.wikipedia.org/w/index.php?title=Contraction_hierarchies&oldid=1173783523 (accessed 11 Mar. 2024).

A

Metaheuristic Solution Methods for MILPs

A.1. Hill Climbing

The hill climbing heuristic is one of the simpler heuristics available, and can be seen as a local search heuristic in its purest form. It starts from an arbitrary solution, possibly found using a greedy algorithm, and then tries to improve the objective function of the problem, until it finds a local optimum. To avoid getting stuck in a local optimum in order to find, or get closer to, the global optimum, one might vary on this method by using several restarts, i.e. starting the algorithm from different initial solutions. The hill climbing heuristic can be used on its own, but more often is simply a part of a larger metaheuristic. For a more thorough discussion on local search heuristics, see for example section 12.2 of Integer Programming [11].

A.2. Variable Neighborhood Search (VNS)

Introduced by Mladenović and Hansen [136], variable neighborhood search is a metaheuristic that is closely related to local search heuristics. The idea of VNS is to systematically change the neighborhood inside of a local search algorithm. To start, several neighborhood structures (i.e. ways of defining a neighborhood) are selected, and an initial solution is found. Every iteration of the algorithm, we cycle through each of these structures. For every structure, a solution is generated at random somewhere in the neighborhood of the previously found solution. Then, some local search method is applied that takes that randomly found solution as its initial solution. If the obtained local optimum is better than the previous solution, this solution becomes the new current solution, defining the end of the iteration. If not, then the found solution is discarded and a new neighborhood structure is searched. The algorithm terminates when none of the neighborhood structures resulted in an improved objective value.

A.3. (Adaptive) Large Neighborhood Search ((A)LNS)

Large neighborhood search is another metaheuristic that resembles local search heuristics, but approaches problems in a slightly different manner. Instead of searching a neighborhood for an improved solution, it uses a destroy method to destruct part of the current solution, and then proceeds to use a repair method to return to a feasible solution. Re-visiting the TSP example, this could mean removing a subset of cities from the route, and adding them back in using a repair heuristic. Typically, the neighborhood that is destroyed gets determined with some element of stochasticity. This ensures that every iteration, the algorithm destroys a different part of the solution. After a destroy and a repair operation, the newly created solution can either be taken as the new solution, or it can be discarded. A simple way of making this choice is discarding the solution if the objective value of the new solution is worse than the previous solution, although this does not need to be the case, and some degree of stochasticity can be introduced here as well. Adaptive large neighborhood search is an extension of LNS. Instead of using a single destroy and repair heuristic for every iteration of the algorithm, ALNS considers a pool of heuristics to choose from. Each of the destroy and repair heuristics in this pool is attached to a weight, the value of which varies based on success of that particular heuristic. This allows ALNS to reach a more varied set of neighborhoods that can be reconstructed in a way that was empirically proven to be most effective. For more information, refer to the chapter written on this topic by Pisinger and Røpke [137].

A.4. Tabu Search

Introduced by Glover [138], tabu search uses a more sophisticated manner of attempting to find a global optimum over a local one. When a local optimum is found, this algorithm moves to a solution in the neighborhood that has inferior objective value, but may lead to an improved local, or even global, optimum later on. If applied without caution, this method will likely start cycling, and we will see the same solutions repeatedly. To avoid this, a tabu-list is created and updated at every iteration, that keeps track of recently found solutions that are forbidden ('tabu') to be visited again. The tabu list does not need to be very large to successfully escape local optima, and can thus be an efficient method for this purpose.

A.5. Simulated Annealing

Unlike the heuristics used before, simulated annealing is a probabilistic technique that randomly chooses what solution to move to at every step. If a solution has a better objective value than the current solution, it will move to that solution with probability 1. If the objective value of a proposed solution is worse than the current value, the algorithm will move to that solution with some probability between 0 and 1 that is reflective of how much worse the proposed solution is. To ensure that the algorithm converges in the long run, the probability that a worse solution gets chosen slowly decreases the more iterations have been done. Section 12.3 of Integer Programming [11] discusses this topic in more detail.

A.6. Genetic Algorithms

Genetic algorithms are slightly more complex than the previously mentioned methods, and form a subcategory of methods called evolutionary algorithms. The idea is the following: instead of starting from a single solution that updates every iterations, genetic algorithms use a set of solutions that are called a generation. Every iteration, the fitness (i.e. the objective value) of each individual of the generation gets assessed. The stronger individuals are, the more likely they are to impact the next generation. The next generation is created using three rules: selection rules (what individuals are used to contribute to the next generation), crossover rules (how to combine parent solutions to create child solutions) and mutation rules (how to mutate a parent solution to create another child solution). More information can for example be found in section 12.3 of Integer Programming [11].

A.7. Ant Colony Optimization

Another more complicated probabilistic technique is ant colony optimization. This method was inspired by the behavior of ants finding routes between their nest and food sources, that find the shortest path by preferring to follow trail pheromones that were deposited by ants that previously made the trip. Before this algorithm can be applied, the problem that needs to be solved needs to be converted to a shortest path problem. It can be applied directly to problems already of this nature, such as the TSP, and is commonly used for routing based problems. The algorithm starts by sending out artificial ants from different (possibly non-origin) nodes in the graph, that completely arbitrarily try to find paths, by visiting each new location out of the set of unvisited locations with the same probability. While traveling along the nodes, the ants drop pheromones. If a path contains pheromones, an ant is more likely to choose that route the next round. The following iteration, the next horde of artificial ants again try to find a shortest path, but due to the dropped pheromones, they are slightly more likely to take (parts of) the previously found paths. Every iteration, some of the pheromones on the graph evaporate. The longer a route takes, the more time has passed after the ant returns, and the more the pheromones dropped on that path evaporate. When a route is short, the pheromones evaporate less quickly. So, the shorter and thus better routes are increasingly more likely to be used, while longer routes slowly stop being used. Eventually, the route the ants choose will converge to at least a locally, and possibly globally optimal route. For a more in-depth discussion, refer to for example the article by Dorigo [139].

B
VRP Variations

Table B.1 provides an overview of the different VRP variations that have been studied in the litereature. It contains three columns, containing a list of scenario characteristics, problem physical characteristics and information characteristics, and has been based on the taxonomy of Braekers et al. [17]. In this paper, the characteristics that differ from previous taxonomies (such as the taxonomy of Eksioglu et al. [140]) are discussed in detail, and should be referred to (possibly together with previous taxonomies) for more information. A general remark however is that some of these blocks contain different options for a variation, so for example load splitting may or may not be allowed. Other blocks introduce new concept that may or may not apply to a single variation, for example with backhauls. A variation might contain the option to offer backhauls which can be done in two different ways, but it could also not offer this option.

Table B.1: Overview of the different VRP variations

Scenario Characteristics	Problem Physical Characteristics	Information Characteristics
Number of stops on route: - deterministic - partially probabilistic	**Transportation network design**: - directed network - undirected network	**Evolution of information**: - static - partially dynamic
Load Splitting: - allowed - not allowed	**Location of addresses**: - customer on nodes - arc routing instances	**Quality of information**: - deterministic - stochastic - forecast - unknown (real-time)
Customer service demand quantity: - deterministic - stochastic - unknown	**Number of points of origin**: - single origin - multiple origin	**Availability of information**: - local - global
Request times of new customers: - deterministic - stochastic - unknown	**Number of depots**: - single depot - multiple depots	**Processing of information**: - centralized - decentralized
Onsite service / waiting times: - deterministic - dependent - stochastic - unknown	**Time window type**: - restriction on customers - restriction on depot / hubs - restriction on vehicles	**Data characteristics**: - real-world data - synthetic data - both - no data used
Time window structure: - soft time windows - strict time windows - mixed	**Number of vehicles**: - single vehicle - limited number of vehicles - unlimited number of vehicles	
Time horizon: - single period - multi-period	**Capacity consideration**: - capacitated vehicles - uncapacitated vehicles	
Backhauls: - simultaneous pickups and deliveries - either linehaul or backhaul	**Vehicle homogeneity (capacity)**: - similar vehicles - load-specific vehicles - heterogeneous vehicles - customer-specific vehicles	
Node/Arc covering constraints: - precedence and coupling constraints - subset covering allowed - recourse allowed	**Travel time**: - deterministic - function dependent - stochastic - unknown	
	Objective: - travel time dependent - distance dependent - vehicle dependent - function of lateness - implied hazard / risk related - other	

C

Heuristic Solution Methods for the EVRP

C.1. Common Metaheuristic Approaches

To be able to solve large instances of the EVRP, many researchers decided to use metaheuristics instead of exact methods. Popular metaheuristics for the EVRP are (A)LNS, VNS, genetic algorithms and tabu search [5]. Kancharla [38], Goeke and Schneider [46] and Pelletier et al. [50] used ALNS to solve their problem instances. Li et al. [59], Xiao et al. [36] and Lu et al. [37] instead used VNS. Genetic algorithms were used by Futalef et al. [60], Karakatič et al. [118] and Abdallah et al. [110]. Barco et al. [58] instead of genetic algorithm used a related approach and created a differential evolution algorithm. Goeke [141] used granular tabu search as their main solution method. Tabu search most often is combined with ALNS or VNS, a few examples of which are given in the next section. Examples of papers using ant colony optimization are written by Mao et al. [115] and Jia et al. [142].

C.2. Hybrid Metaheuristic Approaches

Due to the difficulty of the problem, researchers often use more than a single solution technique. A common strategy is to split the algorithm up in a routing and a charging phase. Froger et al. [39], Montoya et al. [100] and Koç et al. [119] for example solve this by considering the Fixed Route Vehicle Charging Problem (FRVCP) as a subproblem of their respective EVRP variant, in combination with using an iterative search heuristic. The FVRCP determines the charging operations necessary to complete the route that was assigned to a vehicle. This subproblem can be solved both exactly and with a (greedy) heuristic. In general, whenever charging stations are inserted into a solution it can be done heuristically or optimally. In addition to solving this subproblem exactly, another optimal way of doing this optimally is using a labeling algorithm based on dynamic programming. This algorithm has essentially the same goal as solving the FVRCP, providing an optimal set of charging station insertions for a given route. This approach is used by Hiermann et al. [56], Küçükoğlu et al. [113] and Roberti and Wen [54]. Hiermann et al. implemented this method in combination with ALNS, Küçükoğlu et al. combined it with their hybrid simulated annealing/tabu search algorithm and Roberti and Wen used it as the final step in a general VNS heuristic. Other combinations of metaheuristics are for example used by Zhang et al. [63] that enhanced their ALNS algorithm with fuzzy simulation, Li-ying and Yuan-bin [143] that combined ALNS with tabu search, and Ding et al. [144] that combined VNS with tabu search.

C.3. Alternative Approaches

The methods that are given earlier are not the only ways in which one can solve such a problem; plenty of researchers have gotten creative with their solution methods. Some researchers combined ideas from known metaheuristics with other techniques. This applies for example to Li et al. [111], that designed a mixed algorithm containing tabu-search, and Erdoğdu and Karabulut [109] that combined simulated annealing with a constructive heuristic (i.e. a heuristic that starts with an empty solution, which it repeatedly extends until a feasible solution is found) and a local search heuristic. Zhao and Lu [57] created a heuristic approach based on ALNS which contains Integer Programming as a way of improving the heuristically found solution. For a larger overview of the used metaheuristics and the alterations for solving variations of the EVRP, refer to the literature review of Küçükoğlu et al. [5]. Other researchers used entirely different methods. Yao et al. [117] used Benders Decomposition [145] as a way of splitting the problem into a suband master problem, while Yang et al. [55] used the Alternating Direction Method of Multipliers [146]. Yet another possible approach is to define a two-part optimization scheme (often but not necessarily a routing and charging stage) that splits the problem into two (MI)LPs, which then get solved iteratively using a commercial solver or an algorithm such as Bellman-Ford [147]. This is done for example by Chen et al. [116], Yao et al. [148], Basso et al. [51] in 2019 and Basso et al. [62] in 2021. A final approach we would like to mention is the work by Schoenberg and Dressler [149], which used a multi-criterion shortest-path search algorithm using contraction hierarchies [150] to apply an adaptive charging and routing strategy to solve instances of their model.