Erasmus University Rotterdam

Erasmus School of Economics

Master Thesis Analytics And Operations Research In Logistics

Beyond Efficiency: Exploring the Cost Effects of Prioritizing Driver Happiness in Vehicle Routing

Tom Bruinink (504523)

Supervisor:	Twan Dollevoet
Second assessor:	Bart van Rossum
Supervisor ORTEC:	Lotte Berghman
Date final version:	October 20, 2023

Abstract

The standard objective in Vehicle Routing Problems (VRP) both in general and in the ORTEC routing software is cost minimization. However, due to a shortage of drivers and a stronger focus on employee well-being, there is an increasing need to take driver satisfaction explicitly into account when planning routes. We introduce the concept of "driver happiness" by measuring workload balance and region consistency. We evaluate both a two-step approach where a driver assignment is done after solving the standard VRP and an integrated approach where driver happiness becomes a second objective. Based on a case study from Pinnacle Mart and using the ORTEC routing software, we have identified opportunities for significant improvements. However, focusing exclusively on one measure often results in neglecting the other. To address this, we propose combining workload balance and region consistency to enhance overall driver happiness and obtain a wide range of solutions. This flexibility empowers individuals to prioritize and choose the balance that best aligns with their preferences.

Contents

1 Introduction

In the logistics industry, efficient vehicle routing is crucial for timely delivery of goods and services. Traditional approaches to solving the Vehicle Routing Problem (VRP) have primarily focused on cost minimization and CO2 emission reduction. However, these company-centric strategies have proven to be insufficient in meeting the demands of developed markets, where competition is fierce and drivers have increasingly high needs. To adapt to this changing landscape, a more "person-first" approach is necessary, one that considers the well-being and satisfaction of drivers.

Driver shortage has become a pressing challenge in the e-commerce sector, leading to delays and inefficiencies in last-mile delivery. In the USA, the average turnover rate for drivers, indicating the yearly rate at which employees leave a company and are replaced by new hires, has exceeded 100% according to a study by American Trucking Associations (2007). In comparison, the turnover rate for all jobs was reported as 23.7% by Bureau of Labor Statistics (2007), highlighting a notable disparity. The total cost of a turnover can easily reach $15,000 (Min & Lambert, 2002). Recognizing the need to address this issue, researchers and practitioners are exploring new avenues to optimize vehicle routes by incorporating driver happiness as a critical factor. By taking into account drivers' preferences, it is possible to improve their job satisfaction and, consequently, enhance the overall efficiency of the delivery process.

This thesis aims to contribute to the evolving field of vehicle routing optimization by introducing a driver happiness measure that quantifies driver satisfaction based on multiple considerations. Factors such as familiarity with specific regions and balance of workload will be taken into account to construct a measure that assigns a happiness level to each driver for every possible route. The integration of this happiness measure into the optimization process will enable the selection of routes that not only minimize costs or CO2 emissions but also prioritize driver satisfaction. This approach aims for a balance between cost and driver well-being in vehicle routing.

This shift towards prioritizing driver satisfaction acknowledges the changing dynamics of the market and the need to address the driver shortage challenge. Through this research, valuable insights and practical solutions are provided to bridge the gap between traditional cost-focused strategies and the emerging importance of driver well-being in the optimization of vehicle routing.

In the first and second approach presented in this thesis, routes will be designed based on cost minimization. In the first approach drivers are assigned to routes randomly to create a baseline solution. Next, in the second approach, drivers will be assigned to routes based on their preferences.

The third approach will be an integrated approach, in which the creation of routes aims to maximize driver happiness while ensuring that the total cost of all routes remains considerably low. This approach entails balancing driver preferences and costs during route creation, with the goal of optimizing driver satisfaction.

We find that the random allocation consistently is not able to find efficient solutions. When focusing on workload balance, the driver assignment is already nearly optimal, with only marginal improvements seen with the integrated approach. In terms of region consistency, the driver assignment significantly enhances the random solutions, while the integrated approach outperforms the driver assignment. When both workload balance and region consistency are considered together, the driver assignment and integrated approach strike an effective balance, with the integrated approach excelling in region consistency while maintaining a similar workload balance.

This thesis is structured as follows. Section 2 presents the existing literature related to this thesis. In Section 3 the literature is verified by interviews. Section 4 discusses the problem and assumptions. Next, the different solving techniques are presented in Section 5. Section 6 discusses the case study. Section 7 presents the results. Finally, Section 8 and Section 9 discuss the outcomes and conclude.

2 Literature Review

In this literature review, we first focus on understanding driver happiness in detail. Then we discuss the different approaches in research to improve driver happiness.

2.1 Driver happiness

According to a study by Kozyreff et al. (2022), driver happiness can be quantified based on different factors. One key factor is creating similar routes, which gives drivers consistent assignments aligning with their experience and thereby increases their happiness. Another factor is scheduling working hours that closely match drivers' preferred hours. Additionally, taking into account the number of deliveries could be important, as some drivers may prefer more deliveries while others prefer longer driving distances with fewer customers.

We expand upon the concepts presented by Kozyreff et al. (2022). Similar to their approach, we assess driver happiness by examining the alignment of working hours, which we interpret as workload balancing, and consistency, which can be evaluated in various ways. This sections reviews previous studies on workload balancing and discusses research for achieving consistency. Finally, our review concludes by combining the concepts of workload balance and consistency.

2.2 Workload balance

Workload balance is extensively studied in various fields. In the context of workforce scheduling, which is applied in domains like healthcare and public transport, the focus of workload balance is primarily on achieving fairness (Martin et al., 2013; Uhde et al., 2020). This entails evenly distributing less popular shifts, such as Wednesday and Friday evenings, among the staff members (Wynendaele et al., 2021). Nevertheless, it is important to consider the trade-off between fairness and attractiveness, as slightly loosening the fairness requirements can lead to significant improvements in attractiveness or quality of the rosters (Breugem et al., 2022). Research from Van Rossum et al. (2023) supports this, showing that the most cost-efficient solutions often lead to unfair assignments, while minor efficiency sacrifices yield significantly fairer results.

According to Matl et al. (2019), there exist multiple measures for workload balance in transportation problems. The most common measure utilized for achieving workload balance is tour length, which can be quantified in terms of time or distance. However, other measures are more appropriate in specific sectors. For instance, the number of stops is a relevant measure for small package delivery, service time is significant in technician routing and home healthcare services, and load/demand is evenly important as tour length in grocery distribution (Matl et al., 2019).

In their study, Matl et al. (2017) propose multiple approaches to quantify workload balance. These include minimizing the maximum workload (min-max), its lexicographic extension, measuring the range of different workloads, the mean absolute deviation, the standard deviation, or the Gini coefficient. When workload is determined by a variable-sum metric like tour length, numerical optimization typically favors monotonic measures of inequality, such as min-max or its lexicographic extension. Furthermore, according to the findings of Matl et al. (2019), incorporating more complex equity functions such as lexicographic, mean absolute deviation, standard deviation, and Gini coefficient can offer a wider range of potential solutions. Similarly, balancing distance offers more solutions than load, and load offers more solutions than the number of stops. Having a larger set of solutions increases the likelihood of finding a solution that closely aligns with the desired balance outcomes.

In this thesis, we will assess workload balance using tour length, which represents the duration of routes. To build upon the measures mentioned in the literature and integrate them into our research, we will quantify the degree of workload balance by minimizing the standard deviation. This metric assigns greater weight to higher deviations, providing the model with an extra incentive to generate routes that have comparable workloads.

2.3 Consistency

Consistency in service delivery has been identified as another important factor in improving driver happiness. Several studies have explored the concept of consistency in different directions.

Yao et al. (2021) propose the Consistent Vehicle Routing Problem (ConVRP), which aims to create consistent paths by discounting the cost of roads that are traversed every day. Groër et al. (2009) suggests that consistent service can be achieved by prioritizing customers who require service on multiple days and arranging their visits in a specific order. Kovacs et al. (2014) proposes three directions for achieving consistency: arrival time consistency, delivery consistency and person-oriented consistency. Arrival time consistency can be ensured by scheduling customer visits at consistent times throughout the day in the long term. Delivery consistency can be achieved by replenishing customer stocks at regular intervals with consistent quantities or by maintaining stable inventory levels. Person-oriented consistency can be examined from two angles: the customer's viewpoint and the driver's perspective. Driver consistency, where a customer is visited by the same driver all the time, enhances the service quality for the customer. This approach strengthens the customer-driver relationship and allows for personalized service, ensuring person-oriented consistency from the customer's viewpoint. Achieving person-oriented consistency from the driver's perspective entails assigning each driver to the same region or even to the same set of customers. These concepts are known as region consistency and customer consistency, respectively. Kalcsics (2015) proposed the idea of improving region consistency by designing districts. Districting refers to the task of grouping small geographic areas into larger clusters known as districts. The objective is to create districts that are balanced, contiguous, and compact in their geographic layout. Kant et al. (2008) suggest a method for allocating drivers to predetermined routes. This approach involves identifying an anchor point, which represents the center of each driver's working area. The goal is to minimize the overall cost across all routes. In this context, the cost of a route is calculated as the distance between the driver's anchor point and all the customers on that route. By minimizing this cost, drivers are assigned to regions they are familiar with. Additionally, Bender et al. (2020) present a solution approach capable of designing districts that achieve a high degree of operational feasibility and workload balance. Kovacs et al. (2015) found that increased driver consistency and arrival time consistency can be achieved with only a slight increase in travel time.

In this thesis, our main focus will be on region consistency, which emphasizes the alignment of routes with the perspective of the driver, considering their preferences and experiences. This approach is supported by the research conducted by Kozyreff et al. (2022), which suggests that increasing the similarity of routes can enhance driver happiness. In line with the research by Kalcsics (2015), our approach to design regions will involve grouping smaller geographic areas into larger clusters. These clusters will form the basis for our regions. Similar to the approach in Kant et al. (2008), the regions are not intended with strict area assignment, where all customers on a route belong to the same region. Instead, drivers will have the flexibility to select the regions they prefer to serve based on their individual preferences. Consequently, the emphasis of region consistency will be on enhancing driver happiness by assigning them to regions they prefer, rather than solely focusing on efficiency through assigning drivers to familiar regions.

2.4 Combining workload balance and region consistency

Janssens et al. (2015) and Kozyreff et al. (2022) propose a multi-objective optimization problem aimed at minimizing total costs, deviations from region assignments in the tactical planning (region consistency), and workload imbalances. The approach of Janssens et al. (2015) involves creating a tactical planning framework that incorporates microzones assigned to vehicles. In the operational plan, microzones can be reallocated to different vehicles with the aim of achieving balanced and robust assignments. The objective is to minimize differences between microzone assignments in the tactical plan and the operational one, while also reducing total travel distance and workload imbalances. Workload imbalances are assessed by comparing the operational workload with the average workload. In contrast, Kozyreff et al. (2022) aims to minimize the disparity between a driver's familiar route, considered the reference route, and the route assigned for the day. Workload imbalance is measured as the variance between the driver's working hours and their preferred working hours. These approaches ensure a balance between minimizing costs, achieving workload balance, and maintaining region consistency. The authors emphasize the importance of maintaining a healthy balance among the three objective functions throughout the optimization process.

In this thesis, the workload balance and region consistency aspects will first be considered independently and subsequently be combined. In the third approach, known as the integrated approach, the total costs will be evaluated using a weighted objective, inspired by the concept introduced by Janssens et al. (2015) and Kozyreff et al. (2022). This objective encompasses both routing costs and driver happiness costs — either reflecting workload balance, region consistency or a combination of both. However, the workload will be balanced over a longer period, differing from the scope of Janssens et al. (2015) and Kozyreff et al. (2022) where it is addressed on a daily basis. Kozyreff et al. (2022) acknowledge that a longer period may be more suitable for balancing workload.

3 Interviews

To examine whether the preferences of drivers align with the statements made in the existing literature, we conducted interviews with drivers. We began with open-ended questions to gain insights into the working schedules of the drivers and to assess the satisfaction with their assigned routes. Subsequently, we asked specific questions related to preferences and fairness. The findings from these interviews will be presented below. The complete set of interview questions and answers can be found in Appendix A.

We conducted interviews with seven drivers from Terracraft Construction. These drivers have fixed regions they normally serve. However, there is some flexibility, allowing them to deliver customers in different districts when necessary. The drivers value having their own regions because it enables them to become familiar with both the customers and the areas. This familiarity helps them in making efficient deliveries at convenient times for their customers.

The drivers prioritize avoiding overtime but are open to working slightly longer hours on a daily basis. They prefer spreading overtime across multiple days (e.g. adding an hour to their shifts over four days), rather than having a significant amount of overtime on a single day (e.g. working four extra hours in one day). Interestingly, the drivers are not concerned about how much their colleagues work; they primary focus on their own working schedules. Additionally, the nature of the routes they drive (whether they involve long travel times with few customers or short travel times with many customers) does not significantly influence their preferences. This matches their primary focus on their own working hours and aligns with our choice to prioritize tour length, measured through route durations.

Furthermore, we presented the drivers four different weekly scheduling options, each corresponding to different measures: maximizing minimum working time, minimizing maximum working time, range, and standard deviation. Our observations indicate that the drivers consider schedules focusing on range and standard deviation to be the fairest, while they perceive the schedule focusing on maximizing minimum working time as the most unfair. It's worth noting that in the schedule that is considered as the most unfair, none of the drivers work close to the standard 40-hour workweek, with two drivers working 35 hours and two working 45 hours. These results support our decision to prioritize standard deviation in the scheduling process.

,

4 Problem Description

To optimize driver happiness, we propose several measures that quantify driver satisfaction. This thesis focuses on maximizing driver happiness through two main factors: workload balance and region consistency. We start by briefly explaining the basics of the Vehicle Routing Problem (VRP), which traditionally focuses on timely delivery of goods and services. Building on this established framework, we extend the VRP to include the aspect of driver happiness.

4.1 Vehicle Routing Problem

The Vehicle Routing Problem (VRP) is a complex optimization problem focused on finding the best routes for a group of vehicles to efficiently serve a set of customers. The goal is to optimize the allocation of vehicles to routes to meet customer demands, taking into account factors like vehicle capacity, time constraints, and road conditions. Over the past few decades, extensive research has been conducted on the VRP, resulting in numerous variants, such as the Multi-Depot VRP (MDVRP), Periodic VRP (PVRP), and VRP with Time-Windows constraints (VRPTW), among others (Braekers et al., 2016).

In this thesis, customers are considered as tasks that need to be planned. Each task has specific characteristics, including a time window for delivery, a quantity to be delivered, an address, and a duration for completing the task. We assume that each customer is assigned to a department. The departments are independent entities, and customers cannot be served by other departments. Therefore the problem can be decomposed per department. We enforce the constraint that vehicles and drivers are restricted to operating within a single depot. Furthermore we assume that each department consist of a single depot. It is important to note that we make the assumption that both the number of routes and the available drivers are predetermined. Each route is associated with a depot location where it begins after a specified earliest departure time and must be back before a given latest arrival time. Moreover, each route is assigned a capacity, which implicitly determines the vehicle type assigned to it. There are also driver-related constraints to consider. For instance, there is a maximum duration of 8 hours for each route. If a route exceeds this duration, it incurs overtime costs. Furthermore there are regulations regarding breaks for drivers.

The main objective of the VRP is to maximize the number of planned tasks, which represents the number of customers served. A solution is considered feasible when all tasks have been planned. Additionally, the objective is to minimize the total cost. The cost is determined by factors such as the total work time, total distance, and the number of routes employed. It is assumed that the costs, including costs per route, per kilometer, per hour, and per hour of overtime, are uniform for all routes. In this model, the cost per route are 0. This decision arises from our awareness of the predetermined number of routes and the drivers available for the day which implies that the route costs remain fixed and will not play a role in the optimization. It is important to note that the number of planned tasks may vary and is not predetermined or fixed. However, when evaluating different solutions, a higher number of planned tasks is consistently regarded as superior to other alternatives. Subsequently, the solution with the minimum cost is considered the optimal choice.

In the extended version of VRP, an additional dimension is introduced by considering driver happiness. The objective is to maximize driver satisfaction by improving workload balance or region consistency. Therefore the cost objective will be extended to a weighted objective, including both routing costs and costs related to driver happiness. Additionally, we require that the number of unplanned tasks (customers) is zero. To incorporate driver preferences into the model, it is necessary to have knowledge of which drivers will be working. Therefore, we assume that the roster of working drivers is known before the creation of routes.

4.2 Workload balance

We aim to achieve a balanced workload among all drivers over a 30-day period. Our approach aligns with the findings of Nekooghadirli et al. (2022), whose research demonstrates that for planning horizons exceeding five days, solutions tend to be optimally efficient in terms of travel distance and nearly optimal in terms of fairness. However, it is important to note a distinction between our methodology and that of Nekooghadirli et al. (2022). While their approach involves creating all delivery routes in advance and subsequently optimizing the workload balance, our situation necessitates a different approach. Due to tight order deadlines, we are constrained to create routes on a daily basis, either on the day of delivery or the day before. Consequently, we have chosen to extend our planning horizon to 30 days, recognizing that finding a nearly optimal solution is more challenging in our case, as we can only consider assignments from previous days alongside today's routes.

Due to daily variations in the number of routes, the corresponding number of drivers also fluctuates.

Consequently, only a subset of drivers is scheduled to work on any given day. To avoid a scenario where drivers who have worked fewer days are consistently assigned longer routes, we consider the average workload over all working days. This effectively eliminates the influence of the number of working days on the assignment of routes.

Randomly assigning drivers to routes on a daily basis can lead to significant variations in workload. To prevent this, we aim to minimize workload imbalances by assigning drivers to routes that closely align with their preferred workload. We assume that all drivers prefer to work an equal amount of time. Therefore, they all have the same preferred workload, which equals the average workload. By minimizing the difference between each driver's actual workload and the average workload, we can increase overall driver satisfaction. The degree of deviation from the average will be evaluated by the standard deviation, since we consider more significant deviations as relatively worse. This means that two individuals working 2 hours more than the average are considered more favorable than one person working 1 hour more and another person working 3 hours more.

4.3 Region consistency

In addition to workload balancing, we address the importance of region consistency in maximizing driver satisfaction. Drivers often travel across different regions as they navigate various routes, and they may have preferences for specific regions. By considering drivers' preferences and assigning them to routes in regions they prefer, we aim to enhance driver happiness. Additionally, this approach results in consistent assignments that align with drivers' experiences, further contributing to improved driver satisfaction (Kozyreff et al., 2022).

A region is defined as a geographically compact area that exhibits a relatively round shape with minimal distortion (Kalcsics, 2015). Each driver indicates their preferred delivery regions. It is assumed that each driver can have one or multiple preference regions, and multiple drivers can share a preference for the same region. Additionally, there may be regions that are not preferred by any driver. Our objective is to maximize, on a daily basis, the similarity between a driver's preference regions and the regions where customers are located on a route. As a result, we do not consider the assignments from the previous days. We calculate a similarity score for each driver on a specific route as the percentage of customers within their preference regions out of the total number of customers on the route. By maximizing the worst-case similarity score on a daily basis, we enhance the overall satisfaction of drivers, ensuring that each driver attains a minimum similarity score every day.

4.4 Combining workload balance and region consistency

Our objective is to achieve the highest level of driver happiness by combining workload balance and region consistency. Our aim is to find a good balance between optimizing both. Similar to the workload balance approach, our target is to achieve a fair balance over a 30-day period. However, we also aim to enhance daily region consistency, ensuring that every driver has a minimal level of satisfaction each day. The assumptions made in the workload balance and region consistency sections remain consistent in this context as well.

5 Methods

This section introduces three different solution approaches, which differ in the ability to improve driver happiness. The first approach is random allocation, where drivers are assigned to routes in a random manner. The second approach, driver assignment, involves assigning routes to drivers based on their preferences. The third approach is the integrated approach where routes are adjusted and created considering the drivers' preferences. The preferences of drivers are measured in terms of workload balance and region consistency.

5.1 Random allocation

The random allocation serves as the initial benchmark solution for assessing driver happiness. Drivers will be randomly assigned to predetermined routes on a daily basis, regardless of their preferences or previous assignments. This approach demonstrates the workload balance and region consistency that result from randomly assigning drivers to routes without taking driver happiness into account.

5.2 Driver assignment

In the driver assignment, the objective is to optimize driver happiness by assigning drivers to predetermined routes. This would enhance driver happiness without incurring additional cost (Kant et al., 2008). Since the routes are not known beforehand, the assignment can only be done based on all previous assignments and today's routes. First we apply the driver assignment using workload balance, subsequently we focus on region consistency and finally we aim to combine workload balance and region consistency.

5.2.1 Workload balance

In the context of workload balance, the objective is to minimize deviations from the average workload. This problem can be solved exactly, providing the best possible outcome for each evaluation method. The approach involves an assignment of longer routes to drivers with lower workload in previous working days and shorter routes to those with higher previous workloads. The key step in this approach is sorting the drivers based on their average workload in descending order and the route durations in ascending order. This aligns with the approach described by Van Rossum et al. (2023), where payoffs in increasing order, are assigned to workers, in decreasing order of their current utility. Operating under the assumption that we only have information about previous assignments and today's route durations, this sorting process allows for the optimal pairing of drivers with routes, thereby reducing workload imbalances across all drivers (Van Rossum et al., 2023). Consequently, this approach guarantees that the deviation from the average workload is minimized on a daily basis for each driver, ensuring the lowest achievable standard deviation every day.

5.2.2 Region consistency

The objective is to guarantee a minimal level of happiness for each driver. This can be achieved by maximizing the similarity score of the driver with the lowest happiness on a daily basis.

We formulate the driver assignment problem as a standard Integer Program (IP) to maximize region consistency. Let R represent the set of routes, D represent the set of drivers, and $s_{d,r}$ represent the similarity score between driver d and route r. The similarity score can be calculated as the percentage of customers within their preference regions out of the total number of customers on the route. This similarity score is independent of history. To solve this problem, we introduce decision variables. We use binary variables $x_{d,r}$ to denote whether driver d is assigned to route r. Continuous variable z denotes the minimum similarity among all drivers. The mathematical model for the driver assignment is as follows:

$$
\begin{aligned}
\max \quad & z && && (1)\\
\text{s.t.} \quad & \sum_{r \in R} x_{d,r} \cdot s_{d,r} \geq z && \forall d \in D && (2)\\
& \sum_{d \in D} x_{d,r} = 1 && \forall r \in R && (3)\\
& \sum_{r \in R} x_{d,r} = 1 && \forall d \in D && (4)\\
& x_{d,r} \in \mathbb{B} && \forall d \in D, r \in R && (5)\\
& z \in \mathbb{R} && && (6)
\end{aligned}
$$

The objective (1) maximizes the region consistency of the least satisfied person, aiming to achieve a guaranteed minimum of happiness to all drivers. Constraints (2) ensure that the region consistency of every driver is higher than or equal to the value z. Constraints (3) and (4) guarantee feasibility. Constraints (3) ensure that every route is assigned to exactly one driver, while constraint (4) ensures that each driver is assigned to exactly one route. Finally, constraints (5) are the binary decision variable restrictions.

Polynomial time proof

We aim to prove the existence of an algorithm capable of solving the driver assignment problem for region consistency within polynomial time. The presence of such a polynomial time algorithm signifies the possibility of efficiently solving the problem, ensuring its scalability to address larger and more complex instances.

The problem at hand can be viewed as a matching problem in a complete bipartite graph. The constraints ensure that each route is assigned to exactly one driver, and each driver is assigned to exactly one route. This corresponds to a one-to-one matching scenario. Our goal is to maximize the minimum region consistency for a driver.

In our complete bipartite graph, each vertex on one side represents a driver and is connected to every vertex on the other side, which represents a route. The algorithm to find the maximum matching in a complete bipartite graph has a time complexity of $\mathcal{O}(V)$, where V is the number of vertices on either side of the bipartite graph. This algorithm runs in linear time and is considered the fastest polynomial-time solution for finding the maximum matching in such a graph.

However, this algorithm focuses on maximizing the cardinality of the matching and does not address the max-min problem. The max-min problem involves finding a matching in a weighted bipartite graph that maximizes the minimum weight over all selected edges. To solve the max-min problem, we can use an adapted version of the Hungarian algorithm. The Hungarian algorithm is an efficient algorithm that solves the assignment problem in bipartite graphs, finding the minimum weight matching. It operates in polynomial time with a complexity of $\mathcal{O}(n^3)$, where n denotes the number of vertices on each side of the bipartite graph.

During the initialization step of the Hungarian algorithm, where n denotes the number of drivers or routes, we need to fill each cell of the similarity matrix with its corresponding similarity score. This initialization step has a time complexity of $\mathcal{O}(n^2)$. The row reduction and column reduction steps are performed alternately until an optimal solution is obtained or the algorithm terminates. Both row and column reduction processes exhibit a time complexity of $\mathcal{O}(n^2)$. Determining whether the current solution is optimal involves, in the worst-case scenario, a time complexity of $\mathcal{O}(n^3)$. This is due to the need to discover augmenting paths within the bipartite graph. Following that, the matrix adjustment step is executed in $\mathcal{O}(n^2)$ time. Since all these processes occur sequentially, summing up their respective time complexities results in a total time complexity of $\mathcal{O}(n^3)$ for the Hungarian algorithm. The Hungarian algorithm is widely recognized as the most efficient polynomial-time algorithm for finding the minimum weight matching in a bipartite graph. However, in our case, we are not interested in finding the matching with the lowest sum of edge weights. Instead, we aim to identify the maximum minimal weight similarity among the edges.

To address the max-min problem, we adapt the Hungarian algorithm by considering the similarity score of each edge individually. We fill the matrix of the Hungarian algorithm with the corresponding similarity scores. We examine whether there exists a matching where all weights of the selected edges are greater than or equal to the similarity score of the chosen edge. To achieve this, we replace all similarity scores that are lower than the weight of the chosen edge by infinity. Subsequently, we apply the Hungarian algorithm to determine if a minimal weight matching with a cost lower than infinity exists. If such a matching is found, it indicates the presence of a matching with at least the similarity score of that chosen edge. We interpret this outcome as a feasible solution and conclude that the optimal objective value of the initial problem is at least equal to the weight of the chosen edge. If a minimal weight matching with a cost of infinity is found, this is perceived as an infeasible solution and we conclude that the optimal objective value of the initial problem is lower than the weight of the chosen edge.

The process starts with sorting all similarity scores in the similarity matrix, arranging them in ascending order within a list. This sorting step has a time complexity of $\mathcal{O}(n^2 log(n^2)) = \mathcal{O}(n^2 \cdot 2log(n)) = \mathcal{O}(n^2 log(n))$. Subsequently, the binary search algorithm is employed to choose the next edge. The binary search initiates with the middle similarity score in the sorted list. If the length of the list is even, we examine the lower middle element. The binary search tests whether the Hungarian algorithm produces a feasible solution. If the middle score yields an infeasible solution, the search is restricted to the left half of the current search space. Conversely, if the middle similarity score yields a feasible solution, the search continues within the right half of the current space. This iterative process of checking for feasible solutions and narrowing the search continues until the current similarity score results in a feasible solution, while the subsequent similarity score leads to an infeasible solution. This current similarity score establishes the maximum minimum similarity score.

The time complexity of the binary search is $\mathcal{O}(log(n^2)) = \mathcal{O}(2log(n)) = \mathcal{O}(log(n))$. Additionally, the time complexity of the Hungarian algorithm itself is $\mathcal{O}(n^3)$, resulting in a combined time complexity of $\mathcal{O}(n^3 log(n))$ for the modified Hungarian algorithm. Note that the sorting process and modified

Hungarian algorithm are executed sequentially and therefore their time complexities can be added. Hence the total time complexity is $\mathcal{O}(n^2 log(n)) + \mathcal{O}(n^3 log(n)) = \mathcal{O}(n^3 log(n))$. An example of this modified Hungarian algorithm will be presented below.

Example 1

In this example, we have 3 drivers and 3 routes. The drivers need to be assigned to exactly one route each. In the table the similarity score of a driver with a route is presented, measured in percentages. The goal is to find the maximum minimum matching using the method described in the proof.

$$\begin{bmatrix} 10 & 20 & 30 \\ 40 & 40 & 60 \\ 40 & 50 & 20 \end{bmatrix}$$

The sorted list is [10 20 20 30 40 40 40 50 60]. We will check whether there exists a feasible solution for a similarity score of 40%, which is the middle score in the sorted list. The modified table looks the following.

$$\begin{bmatrix} \infty & \infty & \infty \\ 40 & 40 & 60 \\ 40 & 50 & \infty \end{bmatrix}$$

The Hungarian algorithm finds a solution with infinite costs. Hence a similarity score of 40% is not feasible. Therefore the search is restricted to the left half of the current search space. Then the middle element has a score of 20%. The modified table looks the following.

$$\begin{bmatrix} \infty & 20 & 30 \\ 40 & 40 & 60 \\ 40 & 50 & 20 \end{bmatrix}$$

The Hungarian algorithm finds a feasible solution, hence a similarity score of 20% is feasible. However we do not know whether this is the maximum minimum similarity score. Therefore we continue with searching in the right half of the current search space. Then the middle element has a score of 30%. The modified table will look the following.

$$\begin{bmatrix} \infty & \infty & 30 \\ 40 & 40 & 60 \\ 40 & 50 & \infty \end{bmatrix}$$

The Hungarian algorithm finds a feasible solution, hence a similarity score of 30% is feasible. Additionally, the next element in the sorted list, with a score of 40%, is already checked and gave an infeasible solution. Hence we found the optimal solution and 30% is the maximum minimum similarity score.

5.2.3 Combining workload balance and region consistency

In this section we aim to combine workload balance and region consistency to optimize driver happiness. We will introduce a weighted multi-objective function that enhances the similarity score of the least content driver on a daily basis while reducing imbalances in workloads. Similar to workload balance we have a planning horizon of 30 days, but we create routes on a daily basis. This approach aims to strike a balance between daily happiness through region consistency and long-term satisfaction through balanced workloads.

For the driver assignment problem, we extend the Integer Program (IP) discussed earlier in Subsection 5.2.2, which focused on ensuring region consistency in the driver assignment problem. Similar to workload balance, we have a planning horizon of 30 days, but we To evaluate workload balance, our approach involves comparing each driver's workload to the average workload, assigning more significant penalties to larger deviations. As a result, our optimization aims to minimize the total squared deviation. We introduce new parameters; $w_{d,r}$ represents the workload of driver d on route r. The workload of a driver considers the average workload of all preceding days and the route duration of route r. The average workload across all drivers working today, taking into account both the preceding workloads of these drivers and today's route durations, is denoted by μ. To quantify the workload balance in the objective, we introduce a continuous variable y. To achieve a desired trade-off between region consistency and workload balance, we incorporate a weight factor λ into the objective function. The mathematical model for the driver assignment is as follows:

$$\begin{aligned} \max \quad & z - \lambda \cdot y && (7) \\ \text{s.t.} \quad & (2), (3), (4), (5), (6) \end{aligned}$$

$$\sum_{r\in R}\sum_{d\in D} x_{d,r}(w_{d,r}-\mu)^2 \leq y \tag{8}$$

$$y \in \mathbb{R} \tag{9}$$

The objective (7) maximizes the driver happiness. This happiness score is determined by the difference between the minimum similarity score among all drivers z and a product of a weight factor λ and the workload balance value y. Constraint (8) ensures that the total squared workload deviation per day is less than or equal to y.

5.3 Integrated approach

In the integrated approach, the objective is to maximize driver happiness while taking the number of routes as parameter. In this approach, routes are designed to prioritize driver happiness, taking into account the drivers' preferred working hours or preference regions.

We aim to enhance ORTEC's software to prioritize driver happiness. The current software generates routes by looking at multiple objectives. When evaluating different solutions, a higher number of planned tasks is consistently regarded as superior to other alternatives. Subsequently, the solution with the minimum cost is considered the optimal choice. The cost objective in the standard ORTEC software is to minimize the following expression:

$$\text{routing cost} = \alpha \cdot \text{distance} + \beta \cdot \text{working time} + \kappa \cdot \text{overtime} + \eta \cdot \text{routes} \tag{10}$$

The "distance" is the total distance travelled by all vehicles, measured in kilometers. "Working time" is the total amount of regular working hours, and "overtime" represents the number of overtime hours. The variable "routes" is the number of routes on a particular day. The optimal routing cost obtained with the standard ORTEC software is expressed by c^* in the rest of this thesis.

Notably, in this study, we assume that the number of drivers is predetermined. This means the number of routes remains constant, keeping driver costs steady as well. As a result, these driver costs are not considered in the objective. This also makes route-related costs equal 0, setting η to 0. Furthermore, the cost per kilometer (α) is set to 2, the cost per regular working hour (β) is set to 60, and the cost per overtime hour (κ) is set to 75. In the forthcoming sections, this objective will be referred to as the routing costs.

In the extended model, the cost objective will be a weighted objective including both routing costs and costs related to driver happiness. Various weights will correspond to the extent of influence we desire from driver happiness. It is worth noting that greater importance placed on driver happiness is likely to result in higher routing costs. The number of routes is restricted to the number of drivers working on the particular day. It is not possible to create more or fewer routes than the number of drivers. By adopting this approach, we aim to create routes that incorporate the preferences of drivers to improve driver satisfaction.

In general, the optimization in the ORTEC software consists of 3 phases. All operations are performed on the 'current' solution and any time we find a better solution than the current solution, the improved solution is also compared to the best solution found so far.

1. Construction: a solution is built from scratch, trying to plan all tasks. The construction process employs a sequential insertion algorithm to create single routes. This algorithm starts by choosing a seed task, which is typically a difficult task determined by distance. The algorithm then proceeds to plan tasks that are similar or closely related until the route reaches its maximum capacity. Once the sequential insertion algorithm is complete, a local search is performed within the routes. This involves exploring the different neighborhoods, such as two-opt, move and swap operations. These neighborhoods will be further explained later on. The final step in the construction phase is to plan any remaining tasks that could not be initially accommodated by the sequential insertion algorithm.

2. Local search: the solution is improved by exploring similar 'neighboring' solutions. Various commonly used operators, including two-opt, are applied in a configurable order. For such an operator, the quality of the entire neighborhood is estimated, and options are attempted based on this estimation. The search for the neighborhood is performed exhaustively, continuing until no more 'neighboring' solutions with a reasonable estimation, as determined by the current neighborhood operator, can be discovered. The neighborhood operators are cycled through until no further improvements are found. The available neighborhood operators are *two-opt*, *cross-exchange* and *move and swap*.

3. Ruin and Recreate: This is a type of variable neighborhood search in which distant neighborhoods are explored. This is done by 'ruining' the current solution using one of several strategies selected through multiple spins of the roulette wheel. The strategies available in

the roulette wheel include *random cluster removal*, *worst cluster removal*, *trip removal*, and *related removal*. Subsequently, the solution is recreated using the construction and local search strategies until a new local optimum is discovered.

Local search neighborhoods can be divided into two categories: arc exchanges and node exchanges. Arc exchanges involve swapping two arcs within or between routes, while node exchanges involve moving or swapping nodes (customers) to another route or between routes.

1. Two-opt is an arc exchange that swaps two arcs within a route.

2. Cross-exchange is an arc exchange that swaps two arcs of one route with two arcs of another route. This is done in order to swap route segments between routes.

3. Move and swap is a node exchange that can move a group of nodes from one route to another route or swap a group of nodes from one route with a group of nodes of another route.

5.3.1 Workload balance extension

In order to maximize the driver happiness, we aim to minimize workload imbalances among all drivers. We will assess the imbalance by comparing each driver's workload to the average workload. Our approach involves penalizing larger deviations from the average workload more significantly by using quadratic deviations. In our model, we will use different cost settings. The cost setting determines the extra expense for each hour of workload deviation from the average, affecting the cost objective. The cost objective for workload balance is:

$$\min \quad \text{routing costs} + \delta \cdot c^* \cdot TSD \tag{11}$$

$$\text{s.t.} \quad TSD = \sum_{d \in D} (w_d - \mu)^2 \tag{12}$$

The cost per hour deviation is denoted by $\delta \cdot c^*$ and TSD represents the total squared deviation, which is described by the summation underneath. The values of cost per hour deviation are relative to the minimal routing costs and are calculated using the parameter δ, indicating the significance of workload balance in a relative context. For instance, if δ is set at 20%, it means that the total cost increases by 20% of the minimal routing cost for each hour deviation from the average workload. To illustrate, if $\delta = 20\%$ and $c^* = \$1200$, the cost per hour deviation would be \$240. For example, if TSD is 0.5, it leads to an additional cost of \$120. The workload of driver d is w_d, equal to the average workload of all previous days and the current day. The average workload over all drivers working today is represented by μ, taking into account the previous workloads of these drivers and today's route durations. The sum of these squared deviations yields the overall total squared deviation.

5.3.2 Region consistency

In order to maximize the driver happiness among all drivers, we will try to maximize the minimum region consistency. To create a good initial solution, we make adjustments to the construction algorithm. Initially, customers are only assigned to routes that match the preferences of a driver. However, there is a possibility that after creating all routes, numerous tasks may not fit into the preferred routes. In response, for local search and ruin & recreate, we ease the requirement for tasks to be strictly within a driver's preferred region. Within our extended model, we will employ different cost settings. When it comes to region consistency, the cost can be seen of as a reward that is subtracted from the routing costs. The cost objective for region consistency is:

$$\min \quad \text{routing costs} - \gamma \cdot c^* \cdot \text{region consistency} \tag{13}$$

The routing costs are composed of the costs per kilometer, per hour, and per overtime hour, as initially described in the beginning of the section. The reward signifies the reduction in costs corresponding to each additional percentage increase in minimum region consistency. These reward values are expressed relative to the minimal routing costs and are computed using the parameter γ, which represents the importance of region consistency in a relative sense. To illustrate, if γ is set to 0.1%, it implies that the total cost is reduced by 0.1% of the minimal routing cost for each percentage point increase in region consistency. For instance, if $\gamma = 0.1\%$ and $c^* = \$1200$, then the reward amounts to \$1.2 for each percentage point improvement in region consistency. For example, if the region consistency is 15%, this would result in a deduction of \$18 from the total costs.

5.3.3 Combining workload balance and region consistency

In order to maximize overall driver happiness, we aim to combine workload balance and region consistency. The construction will be similar to the region consistency approach, where customers are assigned to routes that align with a driver's preferences. Subsequently, we will integrate the cost objectives introduced in both the workload balance and region consistency extensions. Various configurations for δ and γ will be explored to achieve diverse equilibrium points. This approach provides the freedom to select priorities among various solutions, aiming to strike the good balance between workload distribution and region consistency in alignment with the planners' preferences. The combined cost objective is:

$$\min \quad \text{routing costs} + \delta \cdot c^* \cdot TSD - \gamma \cdot c^* \cdot \text{region consistency} \tag{14}$$

6 Case study

We analyze the performance of our methods using data from a retailer: Pinnacle Mart. This company sells goods to customers and wants to create efficient routes for timely delivery of the goods. In this section, we will discuss the details of our specific problem instances and examine their characteristics. We use distinct sets of instances for different evaluation measures. For region consistency, we analyze 75 instances, each representing a single day's planning. These instances can be categorized into three groups, each containing 25 cases. The 'small' instances involve up to 5 routes, 'medium' instances consist of 7 to 10 routes, and 'large' instances encompass a minimum of 12 routes. In the cases of workload balance and the combination of workload balance and region consistency, we extend our planning horizon to 30 days. In these scenarios, each instance corresponds to a 30-day schedule at a single depot. We examine 10 instances, each located at a different depot. Furthermore, we restrict our analysis to instances that have no unplanned tasks. When evaluating region consistency, we only consider instances without unplanned tasks. For workload balance and the combined evaluation, we exclusively focus on instances that do not contain unplanned tasks for the initial cost setting in the integrated approach. This is because all planning decisions are influenced by the previous day's assignments, making it practically infeasible to filter out all plans with unplanned tasks for all different cost settings. As workload balance and the combined approach operate under distinct initial cost settings, there are slight differences in the instances under consideration.

6.1 Workload balance

In our case study, the workload is balanced over a planning horizon of 30 days, representing a one-month period. Workload is measured as worked time in hours, represented by the route durations. As described in the problem statement, drivers may not work every day, and the number of drivers working each day may vary. The drivers working on a day is predetermined and independent of previous assignments. The workload is balanced daily by considering both the previous workloads and the route durations for the current day. This approach allows for effective workload distribution and optimization throughout the month.

In Table 1 we present the descriptive statistics for workload balance. We consider 10 problem instances, each representing a 30-day schedule at a single depot. Below, you will find the overall statistics including the average, standard deviation, minimum, and maximum values derived from these instances. An interesting observation is that, on average, there are approximately 6.1 drivers working on any given day, while the total number of available drivers at the depot averages around 12.2. This indicates that drivers work half of the days within the scheduling period. Interestingly both the number of working drivers on a day and the total number of drivers exhibit significant variability, with a minimum of 4.9 and 9.2 respectively, and a maximum of 9 and 20 respectively. Moreover, there is a substantial diversity in capacity per vehicle, ranging from 90 to 200. Additionally, it is worth noting that several key parameters, such as order amount per customer, time-window length, and maximum worktime and overtime, are very close across the different instances.

Table 1. Descriptive Statistics workload balance

	Average	Standard deviation	Minimum	Maximum
Number of drivers working	6.1	1.2	4.9	9.2
Total number of drivers	12.2	3.1	9	20
Number of tasks	67.2	14.7	48.1	104.3
Capacity per vehicle	152.9	48.3	90	200
Order amount per customer	6.1	0.3	5.7	6.8
Time-window length (h)	2.8	0.2	2.3	3.0
Maximum worktime (h)	8.1	0.6	6.7	9.0
Maximum overtime (h)	1.3	0.4	0.4	1.5

Note. The descriptive statistics contains 10 instances with a planning horizon of 30 days.

6.2 Region consistency

In our case study, we focus on optimizing region consistency on a daily basis, aiming to maximize the minimum consistency. As a result, our analysis considers a time horizon of a single day. Consequently, we do not take previous assignments into account, as our primary concern is the daily optimization of region consistency.

A region is a geographical area defined by three 5-digit zip codes, grouped based on the distance between them. The creation of regions is based on a distance matrix that captures the distances between the center points of each zip code. The number of regions per depot varies, ranging from 9 to 82, with the majority falling between 10 and 25. To provide flexibility for drivers, each driver has the option to select a certain number of preference regions. The number of preference regions a driver can choose depends on the total number of regions of a depot. For depots with 15 or fewer regions, drivers can choose three preference regions. For depots with 16 to 25 regions, drivers are allowed to choose four preference regions. For depots with more than 25 regions, drivers can select five preference regions. These preference regions can be located in completely different directions from the depot.

Table 2 shows the descriptive statistics for region consistency. We have divided the problem instances into three groups, each containing 25 instances. Small instances have up to 5 routes, medium instances have 7 to 10 routes, and large instances have at least 12 routes. A total of 75 instances will be studied. It is interesting to observe that the Zipcode/Driver ratio, which represents the number of distinct zipcodes visited in a day divided by the number of drivers on a day, is more or less similar for the different instance sizes. This indicates that the number of zipcodes per driver does not depend on the instance size. Notably, in smaller instances, there is an average of 3 drivers working on any given day. This situation might pose challenges in maintaining high region consistency. The limited number of drivers means that tasks can only be distributed among these 3 individuals, increasing the likelihood of numerous orders falling outside the preferred regions of any driver.

Table 2. Descriptive Statistics region consistency

	Average	Small instances	Medium instances	Large instances
Number of zipcodes	32.1	13.1	32.9	50.3
Number of drivers working	8.1	3	7.7	13.6
Zipcode/Driver ratio	3.9	3.9	4.3	3.7
Number of tasks	96.7	28.5	89.9	171.8
Capacity per vehicle	136.8	133.3	132.7	144.5
Order amount per customer	6.1	5.7	5.9	6.5
Time-window length (h)	2.7	2.7	2.7	2.7
Maximum worktime (h)	8.1	8.1	8.0	8.3
Maximum overtime (h)	1.5	1.4	1.5	1.5

Note. The descriptive statistics contains 75 single-day instances, divided in three groups of 25 instances.

6.3 Combining workload balance and region consistency

In the case study, where we combine both workload balance and region consistency, our objective is to enhance driver happiness throughout a 30-day period. Our approach involves optimizing workload balance over this duration, while focusing on optimizing region consistency on a daily basis. This ensures that a minimum consistency score is met each day. As a result, when addressing workload balance, we consider prior assignments, whereas for region consistency, our emphasis lies in daily optimization regardless of previous assignments.

Table 3 shows the descriptive statistics for combining workload balance and region consistency. We consider 10 problem instances, each representing a 30-day schedule at a single depot. An interesting observation is that there is an instance in which the number of drivers working (9.2) significantly exceeds the average (6.3). Having more drivers can potentially simplify the task of creating a well-balanced planning, as there are more available routes to allocate drivers to. Furthermore, we observe another instance with a zipcode-to-depot ratio of 3.1, which is considerably lower than the average ratio of 4.2. This indicates that, on average, drivers have to cover fewer regions. This characteristic might simplify the task of establishing consistent routes for drivers in this depot.

Table 3. Descriptive Statistics combining workload balance and region consistency

	Average	Standard deviation	Minimum	Maximum
Number of zipcodes	26.8	6.5	15.4	41.1
Number of drivers working	6.3	1.1	4.9	9.2
Total number of drivers	12.2	3.1	9	20
Zipcode/Driver ratio	4.2	0.5	3.1	5.1
Number of tasks	69.3	14.8	48.3	104.4
Capacity per vehicle	152.9	48.3	90	200
Order amount per customer	6.1	0.3	5.7	6.7
Time-window length (h)	2.7	0.2	2.3	3.0
Maximum worktime (h)	8.1	0.5	6.8	9
Maximum overtime (h)	1.3	0.4	0.5	1.5

Note. The descriptive statistics contains 10 instances with a planning horizon of 30 days.

7 Results

Results are obtained by applying the different methods explained in Section 5 to the problem instances introduced in Section 6. Firstly, the outcomes for workload balance are showed, followed by the results for region consistency, and finally, the findings concerning the combination of workload balance and region consistency are presented. The provided results present the relation between the incurred costs and the measured driver happiness. Notably, the presented costs refer to the routing expenses, not taking into account the workload balance penalty or the region consistency reward.

As described in Section 6, we require that the planning does not contain unplanned tasks. For region consistency, we only consider results without unplanned tasks. For workload balance and the combined approach, we require that there are no unplanned tasks for the initial cost setting in the integrated approach, since it is practically infeasible to filter out all plans with unplanned tasks for all different cost settings.

The computation time for all instances is under 2 minutes, with the majority completing within 30 seconds. This allows us to test various settings across a large number of instances within reasonable time.

7.1 Workload balance

In this section we compare the performances of the different methods when measuring workload balance. We focus on 10 distinct problem instances. In the context of workload balance, each problem instance entails a 30-day schedule on a single depot, with daily planning optimization. We start by looking at the average results for all instances where any workload deviation from the average workload is penalized. This approach aligns with our ultimate objective of achieving the most equitable solution. After that, we delve deeper into the details.

In Figure 1 and Table 4 the average results for workload balance are displayed. It is evident that the solutions achieved through random allocation are significantly inferior to all other approaches. Notably, the standard deviation in the random allocation approach is 0.462 hours, equivalent to nearly 28 minutes. This implies that such an approach would yield a work schedule where most individuals work between 28 minutes less and 28 minutes more than the average per day. Over a planning horizon of 30 days, this variance translates to a substantial difference of 7 hours less or 7 hours more than average. This assumption is based on our earlier findings in Subsection 6.1, where it was observed that, on average, drivers work 15.0 days.

Zooming in on the driver assignment and integrated approach, we can observe that the driver assignment solution is already very close to optimal, as the integrated approach yields only marginal improvements. In the driver assignment, the standard deviation is 0.072 hours, approximately 4 minutes. Over a 30-day planning period, this deviation amounts to 1 hour, considering an average of 15 working days for drivers. The most significant enhancement achieved by the integrated approach reduces the standard deviation by 23%, equivalent to 0.0168 hours. This reduction translates to a 15-minute difference over a 30-day period, considering that drivers, on average, work 15.0 days. However, it is worth noting that implementing this improvement carries an increase in costs of 5.4%.

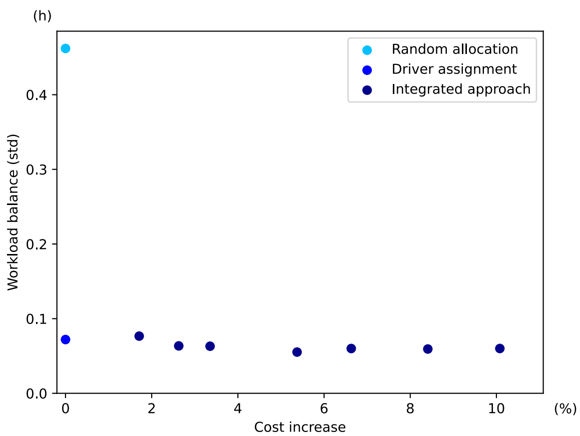

(a) General overview

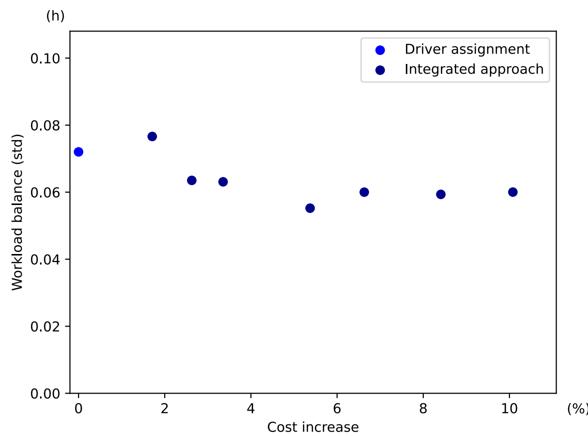

(b) Zoomed in

Figure 1. Average workload balance over all instances

Note. The cost increase are relative to the cost found by cost minimization.

Table 4. Results workload balance

Type	Cost	Standard deviation
Random allocation	2727.34	0.462 hours
Driver assignment	2727.34	0.072 hours
IA, $\delta = 20\%$	2773.97	0.0766 hours
IA, $\delta = 50\%$	2799.00	0.0635 hours
IA, $\delta = 100\%$	2818.83	0.0631 hours
IA, $\delta = 250\%$	2873.89	0.0552 hours
IA, $\delta = 500\%$	2908.19	0.06 hours
IA, $\delta = 1000\%$	2956.66	0.0594 hours
IA, $\delta = 2000\%$	3002.24	0.06 hours

Note. The results represent the mean value derived from 10 individual instances. The term "IA" stands for Integrated Approach. The accompanying percentages δ associated with distinct integrated approach settings indicate the significance of workload balance. For instance, a δ of 20% signifies that, in the context of cost minimization, the total computed cost is enlarged by 20% of the minimal routing cost c^* for each hour deviation for the average.

In the problem description, we assumed that a planning horizon of 30 days is essential to attain a low standard deviation. This differs from the findings of Nekooghadirli et al. (2022), who argued that a planning horizon of 5 days would yield nearly optimal solutions in terms of fairness (workload balance). The standard deviation values corresponding to different planning horizons are presented in Figure 2. This analysis was conducted using the results of the driver assignment, since this approach provided solutions that were near-optimal with minimal costs. It can be observed that there is a sharp decrease in the standard deviation in the initial phase, followed by a lower rate after ten days. However, even at this stage, the standard deviation remains relatively high at 0.27 hours. It is takes 21 days for the standard deviation to drop below 0.10 hours for the first time, and the best standard deviation of 0.068 hours is achieved after 29 days, while the standard deviation after the complete planning horizon of 30 days is very close. This empirical evidence strongly supports the decision to adopt a 30-day planning horizon as it is evidently necessary to achieve the desired low standard deviation.

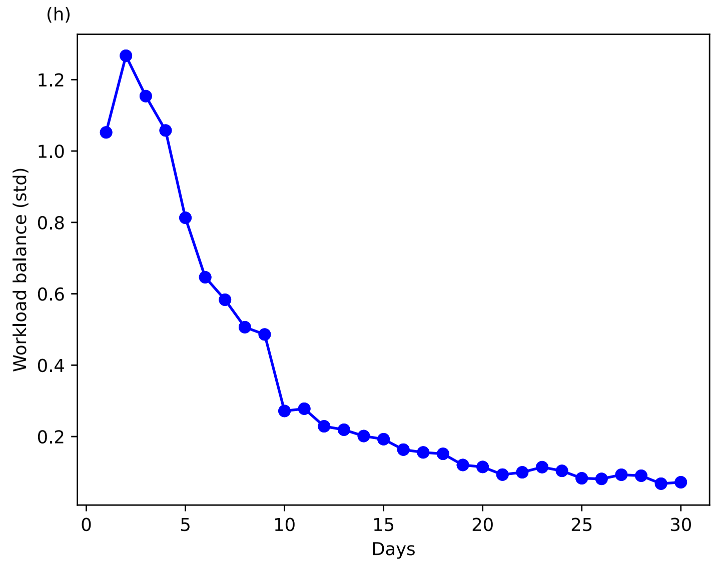

Figure 2. Standard deviation after period of time (derived from driver assignment)

In Figure 3 we examine the standard deviation concerning two planning horizons: one for a week and another for four weeks. When we consider a weekly planning horizon, workloads are reset to zero at the end of each week, essentially starting a new planning at the beginning of each week. The standard deviation is calculated similarly for both planning horizons, involving the measurement of workload deviation from the average workload over a specific time span. We observe that a weekly planning horizon results in a higher standard deviation, particularly after the first week when the workload is reset. At the end of the four-week period, the difference in standard deviation amounts to 0.149 hours. Specifically, the standard deviation for a weekly planning horizon is 0.2396 hours, whereas it is 0.0902 hours for a four-week planning horizon. Over the entire planning horizon of 30 days, this would translate to a difference of 2 hours and 14 minutes. This assumes that drivers work approximately 15.0 days, as discussed in detail in Subsection 6.1. In summary, a planning horizon of four weeks results in a more balanced planning, reducing workload variations compared to a weekly planning horizon.

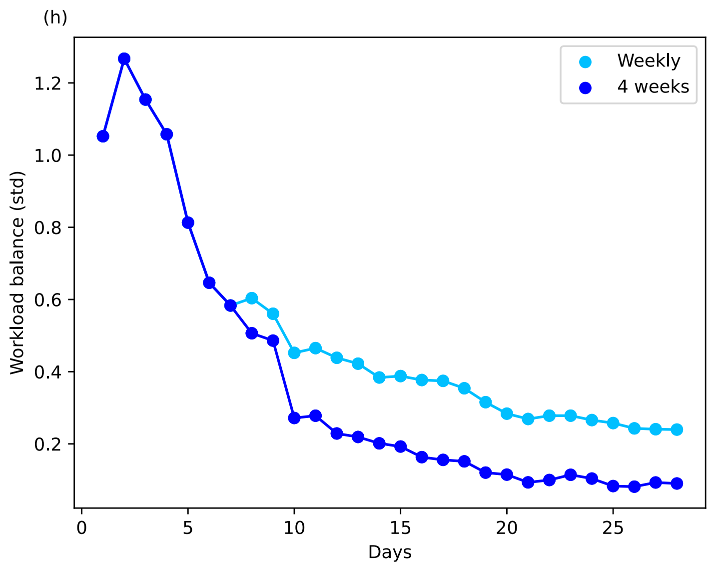

Figure 3. Standard deviation with different planning horizons

7.2 Region consistency

In this section we compare the performances of the different methods when measuring region consistency. We start by examining the average results across all instances, and then delve deeper into various instance sizes and individual schedules.

In Figure 4 and Table 5, we present the average outcomes across all instances for the different methods. It is evident that the random allocation yields a region consistency close to zero across all instance sizes. In contrast, by employing the driver assignment, the region consistency experiences a significant improvement, reaching on average 30.02%. However, the integrated approach outperforms these results, achieving an average region consistency peaking at 43.82% - a substantial 46% enhancement over the driver assignment.
This outcome signifies that, on average, even the least satisfied driver still delivers almost half of their orders within a preferred region. In comparison, the driver assignment method attains less than a third of orders delivered within such regions. It is noteworthy that the average cost increase remains within the 0% to 4% range. Interestingly, a 1.4% cost increase already leads to a region consistency of 40.62%, signifying a 35% improvement compared to the driver assignment. Note that growth in region consistency declines when the value of γ reaches approximately 0.2.

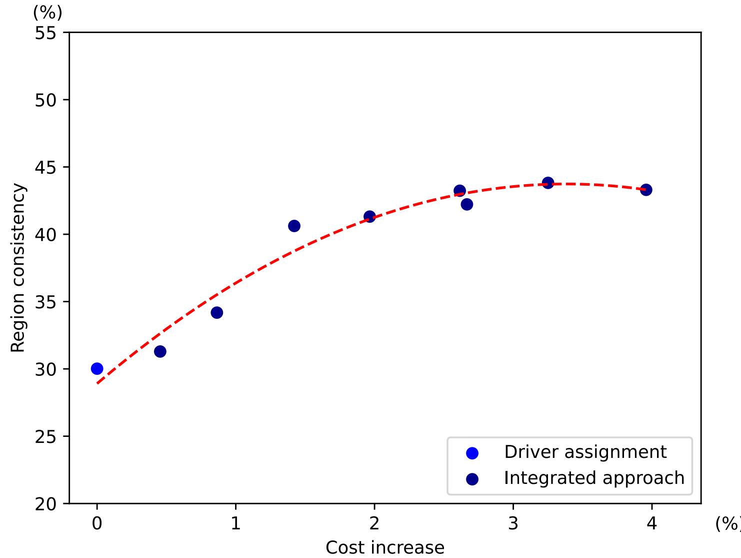

Figure 4. Average region consistency over all instances

Note. The cost increase are relative to the cost found by cost minimization. The trendline is established as a second-degree polynomial.

Table 5. Average results region consistency

Type	Cost	Region consistency
Random allocation	3610.40	4.94%
Driver assignment	3610.40	30.02%
IA, $\gamma = 0.05\%$	3626.83	31.29%
IA, $\gamma = 0.1\%$	3641.59	34.18%
IA, $\gamma = 0.2\%$	3661.72	40.62%
IA, $\gamma = 0.3\%$	3681.36	41.31%
IA, $\gamma = 0.4\%$	3706.66	42.22%
IA, $\gamma = 0.5\%$	3704.80	43.23%
IA, $\gamma = 0.75\%$	3727.77	43.82%
IA, $\gamma = 1.0\%$	3753.30	43.3%

Note. The results represent the mean value derived from 75 individual instances. The term "IA" stands for Integrated Approach. The accompanying percentages γ associated with distinct integrated approach settings indicate the significance of region consistency. For instance, a γ of 0.1% signifies that, in the context of cost minimization, the total computed cost is reduced by 0.1% of the minimal routing cost c^* for every percentage increase in region consistency.

In Figure 5 and Table 6 the region consistency for different instance sizes is presented. Note that for all instance sizes we observe similar trends as for the general results: the random allocation yields a low region consistency, improves considerably with the driver assignment, and the integrated approach outperforms the driver assignment.

It is interesting to note that in small instances, the region consistency achieved by the integrated approach with the lowest weight parameter (γ) is on average a little bit lower than what is attained through the driver assignment strategy. One possible explanation for this phenomenon could be that the emphasis in this integrated approach is predominantly on cost minimization, with insufficient attention given to maintaining region consistency. Furthermore, when dealing with large instances and employing a random allocation method, the region consistency is 0.0% for all instances. This may be due to the presence of a larger number of drivers in these instances, increasing the likelihood that at least one driver is assigned a consistency score of 0.0% when allocated randomly.

Moreover, the region consistency in medium and larger instances significantly surpasses that of smaller instances. For the medium and large instances, the peaks are at respectively 47.1% and 47.02%, while for small instances the peak is at 37.88%. The contrast observed in smaller instances is likely attributed to their lower number of drivers, which leads to a reduced aggregate count of preferred regions shared among all drivers. This outcome could result in a relatively higher portion of customers not falling within any preference region, possibly contributing to the lower region consistency in these scenarios. These findings align with our observations in Subsection 6.2, where we identified that, on average, only 3 drivers operate on any given day.

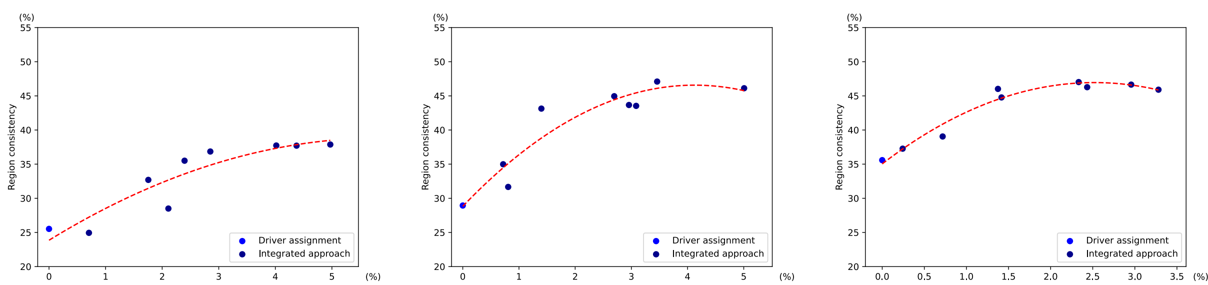

(a) Small instances *(b)* Medium instances *(c)* Large instances

Figure 5. Region consistency for different instance sizes

Note. The trendline is established as a second-degree polynomial.

Table 6. Results region consistency for different instance sizes

Type	Small		Medium		Large	
	Cost	Consistency	Cost	Consistency	Cost	Consistency
Random allocation	1135.48	13.6%	3137.67	1.21%	6558.06	0.0%
Driver assignment	1135.48	25.52%	3137.67	28.94%	6558.06	35.59%
IA, $\gamma = 0.05\%$	1143.49	24.94%	3163.02	31.66%	6573.98	37.27%
IA, $\gamma = 0.1\%$	1159.44	28.5%	3160.22	34.99%	6605.12	39.05%
IA, $\gamma = 0.2\%$	1155.41	32.7%	3181.56	43.14%	6648.20	46.02%
IA, $\gamma = 0.3\%$	1162.68	35.5%	3230.46	43.66%	6650.94	44.77%
IA, $\gamma = 0.4\%$	1167.85	36.85%	3234.47	43.53%	6717.67	46.27%
IA, $\gamma = 0.5\%$	1181.09	37.73%	3222.24	44.95%	6711.06	47.02%
IA, $\gamma = 0.75\%$	1185.16	37.71%	3246.19	47.1%	6751.95	46.65%
IA, $\gamma = 1.0\%$	1191.93	37.88%	3294.77	46.12%	6773.20	45.91%

Note. The results represent the mean values derived from a collection of instances categorized into 25 small instances (up to 5 routes), 25 medium instances (between 7 and 10 routes), and 25 large instances (12 routes or more). The term "IA" stands for Integrated Approach.

7.3 Combining workload balance and region consistency

In this section we compare the performances of the different methods when combining workload balance and region consistency. We focus on analyzing 10 distinct problem instances. Each problem instance involves a 30-day schedule for a single depot, with daily plannings. The average outcomes across all instances are presented in Figure 6 and Table 8. Moreover, we compare different configurations that prioritize specific objectives: cost minimization (via random allocation and driver assignment methods), workload balance, region consistency and different combinations of workload balance and region consistency.

We assign different weights for combining workload balance and region consistency, allowing us to generate a diverse set of solutions. This approach affords planners the flexibility to choose their priorities from a range of solutions, enabling them to achieve the ideal equilibrium between workload balance and region consistency, according to their preferences. The specific weight configurations are detailed in Table 7. For the driver assignment, λ is a parameter used in the objective to set the relative weight for workload balance (y) with respect to region consistency (z) (see (7)). We set λ equal to 50.0. For the integrated approach, we select weights such that the cost increase across different settings is approximately equal. This allows to compare different solutions from the integrated approach without being influenced by cost variations. In the case where the primary focus is on workload balance, we set δ to 250%. This signifies that an hour deviation from the average results in an increase of the total cost by 250% of the minimal routing cost c^*. This selection has previously yielded the lowest standard deviation. Concerning region consistency, we select $\gamma = 1\%$. This signifies that a 1% improvement in region consistency results in a reduction of the total cost by 1% of the minimal routing cost c^*. This choice is made after some experiments as it aligns with the solution with the highest region consistency for small instances and close to the highest score for medium and large instances. For the combined approach, we will assign varying weights to balance workload and region consistency. The values in Table 7 are chosen such that the total costs of the different scenarios are similar.

Table 7. Settings combining workload balance and region consistency

Type	δ	γ
100% WB, 0% RC	250%	0%
75% WB, 25% RC	100%	0.25%
50% WB, 50% RC	50%	0.5%
25% WB, 75% RC	20%	0.75%
0% WB, 100% RC	0%	1.0%

When considering the driver assignment, we observe in Figure 6 and Table 8 that optimizing workload balance does not lead to better region consistency, and vice versa. This implies that improving one aspect of driver happiness does not enhance the other. A similar trend is evident in the integrated approach. However, the combined driver assignment and combined integrated approaches strike a compromise between workload balance and region consistency. Both for workload balance and region consistency, the improvement in objective value when moving from 0% to 25% is much larger than the improvement when moving from 25% to 50% and even more than from 50% to 75% and from 75% to 100%. Interestingly, in all combined integrated approaches, region consistency surpasses the levels achieved when solely focusing on it during driver assignment. This aligns with the findings from the region consistency section, highlighting the potential for significant improvements through the integrated approach. On the other hand, the workload balance achieved using the integrated approach closely resembles that of the driver assignment. When focusing solely on workload balance, the standard deviation is similar for the driver assignment and the integrated approach. Moreover, when combining workload balance and region consistency with equal weights, we observe that the standard deviation for the integrated approach is a little lower than for the driver assignment. This mirrors the results from the workload balance section, indicating that the integrated approach may have limitations in improving workload balance.

In our analysis, we compare the driver assignment to the integrated approach, with our focus on the combined solutions. The integrated approaches show a cost increase ranging from 4.9% to 5.9% when compared to the cost minimization solution (random allocation and driver assignment). We have not included the cost difference between different solutions of the integrated approach in our analysis, as the differences in costs are relatively small. We will compare the integrated approach with equal weights for workload balance and region consistency to the combined driver assignment. In terms of workload balance, the integrated approach reduces the standard deviation by 2 minutes per day (0.1437 hours to 0.1088 hours), totaling a standard deviation of 30 minutes over the entire planning horizon (considering that drivers work for approximately 15 days). Simultaneously, region consistency has improved significantly, increasing from 22.93% to 37.84%. These enhancements highlight improvements in both workload balance and region consistency.

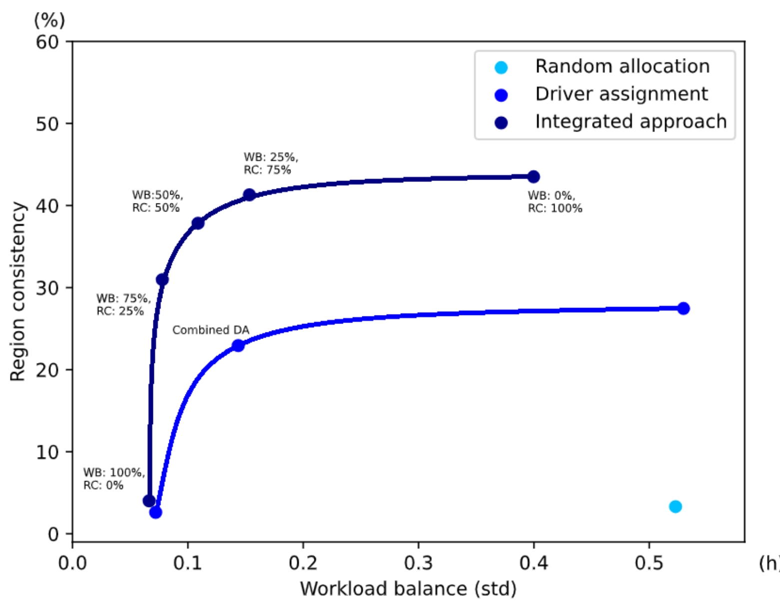

Figure 6. Relation workload balance and region consistency in combined approach

Note. The results represent the mean value derived from 10 individual instances. The labels indicate the corresponding solution. The weights presented are the workload balance and region consistency weights when utilizing the integrated approach.

Table 8. Results combining workload balance and region consistency

Type	Cost	Consistency	Standard deviation
Random	2780.89	3.31%	0.5228 hours
DA, WB	2780.89	2.61%	0.0721 hours
DA, RC	2780.89	27.45%	0.5297 hours
DA, Combined	2780.89	22.93%	0.1437 hours
IA, 100% WB, 0% RC	2943.68	4.02%	0.0664 hours
IA, 75% WB, 25% RC	2915.81	30.97%	0.0778 hours
IA, 50% WB, 50% RC	2925.93	37.84%	0.1088 hours
IA, 25% WB, 75% RC	2926.31	41.3%	0.1533 hours
IA, 0% WB, 100% RC	2916.94	43.5%	0.3997 hours

Note. The results represent the mean value derived from 10 individual instances. The terms "DA" and "IA" stand for Driver Assignment and Integrated Approach, "WB" signifies Workload Balance, "RC" represents Region Consistency.

8 Conclusion

In this thesis, we explored various approaches to integrate driver happiness into the Vehicle Routing Problem, specifically analyzing the cost implications of prioritizing driver satisfaction in vehicle routing. Our primary objective was to develop a strategy aimed at enhancing driver satisfaction by improving workload balance and region consistency.

We quantified workload balance by using the standard deviation of the average workload, serving as a measure of imbalance. Region consistency was defined as the minimum similarity score observed on a single day, with the objective of ensuring a minimum level of satisfaction for all drivers. The similarity score was computed as the percentage of customers within a driver's preferred regions out of the total number of customers on the route. We introduced various approaches for generating routes and assigning them to drivers. Furthermore, we explored various configurations to achieve different trade-offs between driver happiness and costs.

Using data provided by Pinnacle Mart, our initial observations revealed that the random allocation approach is not a suitable method for finding good solutions, characterized by a standard deviation of 28 minutes per day and region consistency near zero. On the other hand, driver assignment emerged as a promising solution strategy. It significantly improved the driver happiness, reducing the average standard deviation to 4 minutes per day and increasing region consistency to 30.02%. When we examined the integrated approach, it outperformed the driver assignment method in terms of region consistency, achieving a notable increase to 43.82% with a cost increase of 3.3%. This indicated that nearly half of the customers served by the least satisfied driver were located within their preference regions. In contrast, with the driver assignment, less than one-third of the customers served by the least satisfied drivers were located within their preference regions. However, concerning workload balance, we found that the integrated approach produced outcomes that were nearly equivalent to those achieved through the driver assignment method. This suggests that the driver assignment method had successfully identified solutions that were close to optimal. In the best configuration, the difference was 15 minutes over a planning horizon of 30 days.

To enhance overall driver happiness, we attempted to combine workload balance and region consistency. At first, prioritizing one measure did not have positive effects on the other. However, by combining workload balance and region consistency, we were able to achieve balances that closely resembled the optimal scores attainable when each measure was individually optimized. We applied different weights to workload balance and region consistency, enabling us to provide a broad range of solutions. This flexibility allows individuals to emphasize their preferences, deciding where to focus more and selecting the solution that aligns best with their specific preferences. The integrated approach improves both workload balance and region consistency at a cost increase of around 5%. Interestingly, our findings align with the trends observed in the workload balance and region consistency methods. The integrated approach offers slightly better workload balance results (30 minutes less standard deviation over the entire planning horizon) compared to the driver assignment. Simultaneously, it shows significant improvements in region consistency, increasing from 22.93% to 37.84%.

In conclusion, this thesis demonstrates the successful integration of driver happiness into the Vehicle Routing Problem. The driver assignment method excels in finding nearly optimal solutions for workload balance, while the integrated approach significantly enhances region consistency and maintains similar workload balance. Moreover, our integrated approach allows us to achieve a good compromise between workload balance and region consistency with a cost increase of 5%.

9 Discussion

The solution strategies presented in this report are able to obtain desirable outcomes, but also have some limitations. To improve the suggested methods, these limitations can be further analyzed and the methods can be extended.

9.1 Possible improvements

In our integrated approach, it is possible that certain tasks remain unplanned due to the heuristic's nature. For region consistency, we restrict to plans where the number of unplanned tasks is zero for all configurations. However, when it comes to workload balance and the combined approach, filtering out plans with unplanned tasks becomes a practical challenge. This is because each plan relies on the workload schedule from the previous day, and unplanned tasks can occur under various cost settings. Consequently, we opted to utilize the first 30 plans without unplanned tasks under the first cost setting as input for all cost settings. This approach allows us to compare routing costs and schedules across different cost settings. In our analysis, we found that there are approximately two unplanned tasks out of a total of around 1800 tasks in a 30-day schedule when using the workload balance approach. To estimate the cost of planning unplanned tasks within the workload balance approach, we added an additional route exclusively for handling unplanned tasks. The average cost associated with each unplanned task was determined to be \$354.40. Consequently, the total cost difference for not planning these 2 tasks over a 30-day period amounted to \$708.80, which is 0.82% of the average total cost of \$86,283.30.

Furthermore, we employ the standard construction algorithm in the ORTEC Routing software, which is a sequential insertion method known for its computational efficiency. For region consistency, it might however be more suitable to consider a parallel insertion algorithm. In the initial phase of the construction heuristic, we prioritize assigning tasks within the preference regions of drivers. When using the sequential insertion method, routes are created sequentially. This sequential process can sometimes result in the last few routes lacking tasks within the preference regions of the drivers. Consequently, these routes end up with a region consistency score of 0. When multiple routes share this lowest consistency score, improving them using local search methods becomes exceptionally challenging. In contrast, a parallel insertion algorithm adds tasks one by one to different routes. This approach significantly increases the likelihood that every route will have at least one task within the preference regions of the driver. Hence, adopting a parallel insertion algorithm can enhance region consistency.

To prevent getting stuck in a local optima when multiple routes share the lowest consistency score, another strategy involves tweaking the local search process. For instance, introducing a penalty for the number of routes with the lowest consistency score can be effective. If improvements are made to any of these routes, the objective value will improve. Another approach involves slight adjustments to the consistency score. The consistency score is defined as the minimum similarity score across all routes. By incorporating the average similarity score of all other routes, enhancements in one route will elevate the consistency score. To maintain focus on the minimum similarity score, assigning a low but still significant weight to the average is essential. This ensures that improvements in the minimum consistency score are always prioritized over enhancements in the average consistency.

In this thesis, a region is considered as a geographical area defined by zip codes, with the requirement of maintaining a minimum distance between them. As noted by Kalcsics (2015), there is a lack of consensus regarding methods for evaluating the compactness, balance, and contiguity of regions. Therefore, additional research aimed at defining what constitutes a "good" region may contribute to the development of more suitable region designs.

9.2 Possible extensions

In this thesis, we assume that each department consists of a single depot, and all customers are served from that depot. As a result, all routes start and finish at the same single depot. However, there are cases where departments are composed of multiple depots. If there are multiple depots within a department, customers can be assigned to routes originating from different depots, allowing for efficient and flexible allocation of resources to meet customer demands. It also means that customers can be assigned to routes originating from different depots on different days. This also means that drivers from different depots might serve the same customer on different days. It would be interesting to investigate multi-depot departments and explore whether they yield better solutions.

Additionally, we make the assumption that all drivers have an equal preferred workload, which means that they prefer to work an equal amount of time. However, in practice, drivers may have different preferences regarding their workload. Additionally, there are variations in contractual hours, with some drivers working full-time while others work part-time. Furthermore, there is a discrepancy in experience levels among drivers, leading to differences in their ability to complete routes efficiently. In terms of region consistency, we make the assumption that all drivers are treated similarly. However, it is important to acknowledge that drivers may vary in their knowledge of specific areas, with more experienced drivers likely having a broader understanding of different regions. Additionally, companies sometimes prioritize assigning routes to drivers with more experience and those who have been with the company for a longer period of time, as a way to recognize and reward their loyalty. Furthermore, there is a tendency for companies to give preference to their full-time drivers, aiming to ensure their satisfaction. Consequently, part-time drivers may often end up being assigned the remaining routes.

In this thesis, the similarity score is determined as the percentage of customers on a route belonging to one of the preference regions of a driver. However, there are various methods for calculating a similarity score. In the study conducted by Smilowitz et al. (2013), the objective is to minimize the number of regions covered by the drivers. A region is considered covered if a driver visits at least one customer in that particular region during the specified time period. Another approach is to define the similarity score based on the number of customers in a route who match the preference regions of a driver. However, this approach may lead to a situation where a region consistency with 3 customers on a route containing 3 customers is considered worse than a consistency with 4 customers on a route consisting of 20 customers.

Moreover, the thesis considers a binary customer-driver similarity: a customer is either in a driver's preference region or not. Exploring alternative approaches could enhance practical applicability. For instance, assigning a score of 0.5 to customers in neighboring regions, giving a negative score to customers in regions that the drivers highly dislike or utilizing distance-based metrics. Kant et al. (2008) present an approach where the total distance of customers to the anchor points, the center points of drivers' working region, is minimized. In our context, the distance for a customer in a driver's preference region would equal 0, while other customers are measured from the nearest preference region center. In that case, minimizing total customer distance from preference regions, rather than maximizing minimum similarity, becomes the objective. This method imposes higher penalties for customers at greater distances, favoring those that are closer. Research should assess whether this significantly enhances driver happiness.

Optimizing the region consistency over a longer period may potentially result in a more equal distribution of region consistency among all drivers. However, achieving this requires optimizing the region consistency of each individual driver, rather than solely focusing on maximizing the minimum consistency. Due to the interdependence of routes and costs, optimizing them in a lexicographic order becomes challenging. Additionally, the approach of balancing region consistency based on the previous assignment may lead to situations where individuals who previously had high consistency may experience significantly lower consistency on a given day if it benefits overall region consistency. Consequently, it fails to prioritize the goal of maximizing the minimum region consistency for all drivers each day. For instance, a driver who initially had 100% region consistency on the first day might experience a drop to 10% on the second day, instead of having 45% region consistency on both days for example. In contrast, our interviews revealed that for workload balance, having a consistently equal workload every day is less crucial, as long as the workload remains relatively consistent over a longer period. The aspect of optimizing region consistency over a longer period is a promising research topic for the future.

References

American Trucking Associations. (2007). *Press release*. Retrieved January 20, 2007, from http://www.trucking.org

Bender, M., Kalcsics, J., & Meyer, A. (2020). Districting for parcel delivery services – a two-stage solution approach and a real-world case study. *Omega*, *96*, 102283.

Braekers, K., Ramaekers, K., & Van Nieuwenhuyse, I. (2016). The vehicle routing problem: State of the art classification and review. *Computers & Industrial Engineering*, *99*, 300–313.

Breugem, T., Dollevoet, T., & Huisman, D. (2022). Is equality always desirable? analyzing the trade-off between fairness and attractiveness in crew rostering. *Management Science*, *68*(4), 2619–2641.

Bureau of Labor Statistics. (2007). *Wages by area and occupation*. Retrieved November 21, 2007, from http://www.bls.gov

Groër, C., Golden, B., & Wasil, E. (2009). The balanced billing cycle vehicle routing problem. *Networks*, *54*(4), 243–254.

Janssens, J., Van den Bergh, J., Sörensen, K., & Cattrysse, D. (2015). Multi-objective microzone-based vehicle routing for courier companies: From tactical to operational planning. *European Journal of Operational Research*, *242*(1), 222–231.

Kalcsics, J. (2015, May). Districting Problems. In G. Laporte, S. Nickel, & F. S. da Gama (Eds.), *Location Science* (pp. 595–622). Springer.

Kant, G., Jacks, M., & Aantjes, C. (2008). Coca-cola enterprises optimizes vehicle routes for efficient product delivery. *Interfaces*, *38*(1), 40–50.

Kovacs, A. A., Golden, B. L., Hartl, R. F., & Parragh, S. N. (2014). Vehicle routing problems in which consistency considerations are important: a survey. *Networks*, *64*(3), 192–213.

Kovacs, A. A., Parragh, S. N., & Hartl, R. F. (2015). The multi-objective generalized consistent vehicle routing problem. *European Journal of Operational Research*, *247*(2), 441–458.

Kozyreff, G., Meerbergen, K., & Zobiri, F. (2022). Driver happiness in e-commerce. Mathematics for Industry, Leuven 2022. Internal Paper.

Martin, S., Ouelhadj, D., Smet, P., Vanden Berghe, G., & Özcan, E. (2013). Cooperative search for fair nurse rosters. *Expert Systems with Applications*, *40*(16), 6674–6683.

Matl, P., Hartl, R. F., & Vidal, T. (2017). Workload Equity in Vehicle Routing Problems: A Survey and Analysis. *Networks*, *52*(2), 239–260.

Matl, P., Hartl, R. F., & Vidal, T. (2019). Workload equity in vehicle routing: The impact of alternative workload resources. *Computers & Operations Research*, *110*, 116–129.

Min, H., & Lambert, T. (2002). Truck driver shortage revisited. *Transportation Journal*, *42*(2), 5–16.

Nekooghadirli, N., Gendreau, M., Potvin, J.-Y., & Vidal, T. (2022). Workload equity in multi-period vehicle routing problems. *arXiv preprint arXiv:2206.14596*.

Smilowitz, K., Nowak, M., & Jiang, T. (2013). Workforce management in periodic delivery operations. *Transportation Science*, *47*(2), 214–230.

Uhde, A., Schlicker, N., Wallach, D. P., & Hassenzahl, M. (2020). Fairness and decision-making in collaborative shift scheduling systems. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–13.

Van Rossum, B., Chen, R., & Lodi, A. (2023). Optimizing fairness over time with homogeneous workers. *23rd Symposium on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems*.

Wynendaele, H., Gemmel, P., Pattyn, E., Myny, D., & Trybou, J. (2021). Systematic review: What is the impact of self-scheduling on the patient, nurse and organization? *Journal of Advanced Nursing*, *77*(1), 47–82.

Yao, Y., Van Woensel, T., Veelenturf, L. P., & Mo, P. (2021). The consistent vehicle routing problem considering path consistency in a road network. *Transportation Research Part B: Methodological*, *153*, 21–44.

Appendix

A Interview questions and results

This section will present the questions and answers from the interviews. For open-ended questions, the most frequently provided answers will be listed. In the case of multiple-choice questions, the frequency of each chosen answer is shown. Additionally, some questions will include comments based on the responses provided by the interviewees. Some multiple choice questions have less answers, as some interviewees argued that they could not give a clear answer to the question.

Interview

1. Hoe zou je een normale werkweek omschrijven? Hoeveel uur maak je elke dag?

- ***Vertrek om 6 uur 's ochtends***
- ***Laden voor volgende dag gebeurt na afloop van de rit***
- ***Aantal uren per chauffeur verschilt: meestal meeste 8/9 uur per dag (40/45 uur per week), sommige uitschieters naar 11 uur (55 uur)***
- ***Aantal uren per dag vrij constant***
- ***15 klanten per route (8/9 als er 2 routes per dag zijn)***

2. Ben je tevreden met de routes die je dagelijks rijdt? En wat zou je willen verbeteren aan die routes?

- ***Chauffeurs zijn over het algemeen heel tevreden over de routes: ze zijn logisch, goed haalbaar en het is fijn dat ze steeds in hetzelfde gebied rijden.***
- ***Chauffeurs passen vaak de routes een beetje aan, op basis van lading, tijdstip dat klant aanwezig is (sommigen zijn er 's ochtends vroeg nog niet), of handiger parkeren/aanrijden (in de stad)***

3. Op welke manier houdt de planner rekening met jouw wensen?

- ***Vast gebied, maar wensen worden niet gevraagd en daar wordt dus niet echt rekening mee gehouden***
- ***Vrij vragen kan prima (voor vakantie, doktersbezoek, kinderen)***

4. Hoe vaak ben je tevreden over de routes die je rijdt voor je werk?

- Nooit ***0 keer***
- Zelden ***0 keer***
- Soms ***1 keer***
- Vaak ***6 keer***
- Altijd ***0 keer***

5. Ik vind het belangrijk om geen overuren of te weinig uren te maken.

- Volledig oneens ***1 keer***
- Oneens ***2 keer***
- Maakt me niet uit ***0 keer***
- Eens ***3 keer***
- Volledig eens ***1 keer***

6. Stel dat je korter of langer dan gepland moet werken. Wat heeft dan je voorkeur?

- X tijd minder ***2 keer***
- Maakt me niet uit ***3 keer***
- X tijd meer ***2 keer***

Comment: Er zijn altijd verschillende werktijden dus maakt niet uit

7. Ik vind het belangrijk dat ik op een dag evenveel werk als mijn collega's.

- Volledig oneens ***0 keer***
- Oneens ***0 keer***
- Maakt me niet uit ***5 keer***
- Eens ***2 keer***
- Volledig eens ***0 keer***

Comment: Ik houd me niet bezig met hoeveel anderen werken, of ze nou 30 of 50 uur werken, dat is aan hen

8. Stel dat er twee routes zijn die tegelijkertijd beginnen en tegelijkertijd eindigen. Rijd je dan liever een route met relatief veel reistijd en weinig klanten of een route met relatief weinig reistijd en veel klanten?

- Veel reistijd en weinig klanten ***1 keer***
- Maakt me niet uit ***4 keer***
- Weinig reistijd en veel klanten ***0 keer***

Comment: Veel reistijd mag geen inefficiënte routes opleveren (midden op de dag op en neer naar een verre klant)

Comment: beide hebben voordelen en nadelen

9. Stel dat je deze week 4 uur moet overwerken. Geef je er de voorkeur aan om op één dag 4 uur langer te werken, of werk je liever 4 dagen lang 1 uur extra?

- Eén dag 4 extra ***0 keer***
- Maakt me niet uit ***2 keer***
- 4 dagen lang 1 uur extra ***5 keer***

10. Ik vind het belangrijker dat ik over een langere termijn mijn contracturen werk dan dat ik elke dag opnieuw mijn contracturen werk.

- Volledig oneens ***1 keer***
- Oneens ***1 keer***
- Maakt me niet uit ***2 keer***
- Eens ***2 keer***
- Volledig eens ***1 keer***

11. Wat zou volgens jou een eerlijk rooster zijn? (Meerdere antwoorden mogelijk)

- Iedereen werkt even lang per dag
- Iedereen bezorgt evenveel lading/m^3 per dag
- Iedereen bezoekt even veel klanten per dag

Comment: Sommige klanten / adressen zijn heel moeilijk (dus evenveel klanten kan niet)
Comment: Vrachtwagens en vormen van lading zijn verschillend dus dezelfde hoeveelheid lading kan niet
Comment: Het maakt de chauffeurs niet uit wat anderen doen, maar meeste lijken gelijke hoeveelheid werk het belangrijkste te vinden.

12. Jij en drie collega's werken deze week en er zijn vier beschikbare schema's. In deze schema's staat hoeveel uur elke chauffeur die week werkt. Je weet van tevoren niet of je chauffeur 1, 2, 3 of 4 bent.

	Schema 1	**Schema 2**	**Schema 3**	**Schema 4**
Chauffeur 1	35u	32u	35u	35u
Chauffeur 2	35u	42u	39u	41u
Chauffeur 3	45u	43u	43u	41u
Chauffeur 4	45u	43u	43u	43u
Minimum	35u	32u	35u	35u
Maximum	45u	43u	43u	43u
Gemiddeld	40u	40u	40u	40u

Welk schema vind jij het eerlijkste?

- Schema 1 ***0 keer***
- Schema 2 ***0 keer***
- Schema 3 ***3 keer***
- Schema 4 ***2 keer***

En welk schema vind je het minst eerlijk?

- Schema 1 ***3 keer***
- Schema 2 ***1 keer***
- Schema 3 ***0 keer***
- Schema 4 ***0 keer***