A supply chain holistic approach to the time window assignment vehicle routing problem with product-dependent variables

André Bronze Ramos Guerra Silva

Masters' Dissertation

Supervisor at FEUP: Prof. Manuel Pina Marques

FEUP FACULDADE DE ENGENHARIA
UNIVERSIDADE DO PORTO

Mestrado em Engenharia e Gestão Industrial

2022-07-11

Abstract

Food retailers face increased pressure from more demanding customers than ever, direct competitors who violently fight for such customers and stressed supply chains. For that reason, in a sector where a small group of companies own hundreds of stores nationwide, it is important to assure that they can deliver all necessary products daily to every store, to avoid stockouts and ensure customer satisfaction. Furthermore, the retailers must guarantee an efficient supply chain and an optimized fleet and continue to search for ways to reduce operational costs.

For all these reasons, it is of insurmountable importance that the retailer can define a delivery schedule that balances client satisfaction with reducing fleet size and operational costs. Such a schedule is based on the delivery time windows of each product in each store.

To do so, this thesis presents an alternative formulation of the Time Window Assignment Vehicle Routing Problem with Product-dependent Variables that intends to obtain a new set of optimized delivery windows, looking at all supply chain costs and not only the travel costs of the original formulation. A three-phase hybrid approach is proposed that starts with a route generation phase that includes a large set of operational constraints to reduce the computational effort of the two-stage stochastic optimization problem formulated for phase two. The last phase is deployed to guarantee the model’s validity.

By applying this new model to a case study of the largest food retailer in Portugal, the preliminary results show that by changing an average of 25% of the existing time windows, a retailer can obtain up to 50% of cost reduction by reducing the required fleet and levelling the number of pallets delivered throughout the day.

Uma abordagem holística à cadeia de abastecimento no problema de roteamento de veículos e atribuição de janelas de entrega com variáveis dependentes do produto

Resumo

Os retalhistas alimentares enfrentam pressões crescentes de clientes cada vez mais exigentes, uma forte concorrência que luta por esses mesmos clientes e cadeias de abastecimento cada vez mais stressadas. Por isso, num setor em que opera um pequeno número de empresas que possuem centenas de lojas espalhadas por todo o país, é importante garantir a entrega diariamente de todos os produtos todas as lojas, evitando ruturas de stock e assegurando a satisfação do cliente. Os retalhistas têm ainda de garantir cadeias de abastecimento eficientes, com recurso a frotas otimizadas e continuar a procurar formas de reduzir os custos logísticos operacionais.

Por todas estas razões, é de enorme importância a capacidade do retalhista definir um plano de entregas que satisfaça o cliente, com uma frota o mais reduzida possível e pelo menor custo possível. Esse plano depende diretamente das janelas de entrega de cada produto, negociadas com cada loja.

Com esse propósito, a presente tese apresenta uma formulação alternativa para o problema de roteamento de veículos e atribuição de janelas de entrega com variáveis dependentes do produto (TWAVRPP) que pretende otimizar as janelas de entrega, reduzindo todos os custos operacionais, e não apenas os custos de viagem da formulação original do referido problema. É proposta uma abordagem híbrida com três fases em que a geração de rotas, na primeira fase, pretende incluir um largo conjunto de restrições operacionais para absorver o esforço computacional do problema estocástico de otimização com dois níveis da segunda fase. A terceira e última fase pretende validar o modelo.

A aplicação deste modelo a um caso de estudo do maior retalhista alimentar a nível nacional permitiu obter um primeiro conjunto de resultados preliminares segundo os quais, ao alterar cerca de 25% das janelas de entrega atuais, se pode poupar cerca de 50% dos custos operacionais do retalhista, através da redução do tamanho da frota necessária e da entrega de paletes mais nivelada ao longo do dia.

Acknowledgements

Firstly, I would like to thank my supervisor Prof. Manuel Pina Marques for all the effort put into the review of this work and for his entire availability and interest shown in helping me achieve the best outcome possible.

Secondly, to Daniela, my supervisor at LTP, I want to thank you for all the insights and feedback provided and for everything you taught me during these four months. Also, I appreciate all the time you spent reading and reviewing my work and all the help in organizing it in the best way possible.

To Horácio and Miguel, the other members of my team, who accompany me in all ups and downs, for all the brainstorming and work we developed together, I thank you for everything I’ve learned with you and for always keeping the bar high. In both of you, I also want to thank all LTPeers for receiving me into this company in the best way I could have asked for and for creating the most fun and exciting environment, where I will never stop learning and never fear asking questions.

To all the professors who have crossed my path, from kindergarten to the university, thanks for being part of this journey. Without you and everything you taught me, I know I wouldn’t be where I am today.

To all of who I have the pleasure of calling my friends, I wanna thank you for all the moments we have spent together. My childhood friends, who grew alongside me and together we got to know the wonders of the world that surround us. My schoolmates, who spent so many hours with me and who are part of so many memories. And my university buddies who faced with me all the challenges of this new adventure that is now ending and that I will cherish for the rest of my life. I don’t know if I wouldn’t be here without you, but I know that it wouldn’t have been as much fun as it was.

Last, but certainly not least, my family, for making me the man I am today, for everything that you have taught me and everything you have invested in me, my deepest and most honest gratitude.

Catarina, my little sister, thank you for being the best sister I could have asked for, even with all the headaches you give me. For always being there, for everything we’ve gone through and for everything we will conquer in the future, I wouldn’t replace you with anyone else.

For my mother, who taught me what true love means, who showed me the meaning of perseverance and courage, and that proved to me that by embracing the things we can’t control we can grow wiser and stronger, I want to thank you for everything that you have always done for me and for ensuring that I would have all the tools at my disposal to be whatever I want to be.

Finally, I have to acknowledge my role model, my father, who taught me that even when facing the most extreme situations, nothing is strong enough to bring us down, when we have the will to fight it. I hope I make you proud every single day, wherever you are.

"It's your road and yours alone.
Others may walk it with you, but no one can walk it for you."
Rumi

Contents

Acronyms

ALNS Adaptive Large Neighbourhood Search

ConVRP Consistent Vehicle Routing Problem

DTWAVRP Discrete Time Window Assignment Vehicle Routing Problem

KPI Key Performance Indicator

LNS Large Neighbourhood Search

LP Linear Programming

MIP Mixed-Integer Program

PBL Pick–by–Line

PBS Pick–by–Store

PVRT Periodic Vehicle Routing Problem

TS Tabu Search

TSP Travelling Salesman Problem

TSPTW Travelling Salesman Problem with Time Windows

TWAVRP Time Window Assignment Vehicle Routing Problem

TWAVRPP Time Window Assignment Vehicle Routing Problem with Product Dependent Variables

VRP Vehicle Routing Problem

VRPTW Vehicle Routing Problem with Time Windows

List of Figures

List of Tables

1 Introduction

1.1 Motivation

The food retail industry is highly dependent on robust supply chain management to ensure that all products get in time from supplier to store, especially when we are talking about highly perishable goods. This industry consists of thousands of producers and even more final customers connected by a small subset of large retailers. A food retailer that manages a large-scale operation is responsible for connecting hundreds of nodes daily, usually with a small fleet of vehicles and the need to move thousands of different SKUs (stock keeping units).

However, beyond this logistic level, another degree of complexity is added when accounting for the company's specifications, all the stakeholders involved, and the environment surrounding them. If in a perfect world, the best supply chain would allow any product to be transported at any time by any vehicle between any nodes of the chain, in the real world, that would be impossible due to all the constraints that these operations are subject to. For that reason, a planner needs to account for the costs of the entire operation, the timings of delivery, all the travelled distances, the country's laws regulating work and food quality, the physical characteristics of the nodes involved, and much more.

Another reason these retailers must bet on effective supply chain management is the competition they face. According to “Portugal Grocery Retailer Preference Index: Second Edition” (2020) by Dunnhumby, the self-acclaimed global leader in Customer Data Science, the food retail sector in Portugal “is dominated by a small number of large Retailers that specialise in traditional offerings. Large stores with broad assortments, well subscribed loyalty schemes, and high-low pricing models make up most of the market here”. And according to the second edition of this report, the “One-Stop-Shop” pillar has become the most important to the Portuguese consumer incrementing, even more, the necessity to ensure that all products are always available at every store.

1.2 Project Description

The present thesis was developed in an Advanced Analytics & Business Consultancy company to create a model that allowed for the optimization of product-dependent time windows, but with an inclusive look into all logistical costs of this kind of operations. For this, it a partnership was established with the largest food retailer operating in Portugal to obtain a real-world case study to validate the results. The main sponsor of this project was the retailer’s Logistics Department, which is responsible for the daily operation of delivering products from the company depots to the stores. In the edible category, this operation consists of two large warehouses that receive the products from the suppliers and send them, daily, to stores around the country. In the last few months, this company faced a set of new challenges:

1) The incorporation of a new store concept into the logistic operations brought a new and large set of small stores spread across the country.

2) A set of new constraints related to changes made to the warehouses and new logistic flows.

3) A new routing software that allows new tests and validations.

All these challenges, when associated with the recurring problem of an ad-hoc definition of the time windows used to deliver the products to each store, complicated the management of the current fleet. On top of this, the impact of COVID on the operations made clear a necessity to have a new set of windows more flexible for large and unpredictable fluctuation of demand that would not create bottlenecks at the warehouse loading operations or create the necessity of renting more vehicles.

This junction of factors made clear the necessity of a new project to optimize the delivery time windows for each store, incrementing the fleet's utilization rate and minimizing the entire process's logistic costs.

1.3 The main objectives of the project

To obtain a set of new and optimized delivery time windows to each store, it was important to have a solid schedule of activities and a structured plan of what to do. During the planning phase, three main goals were defined for the execution phase. The first one was to map all logistic flows of the operation, collect all important data, access all the constraints that all main stakeholders were subject to, and identify the cost drivers involved with the problem. The second objective was to create an analytical model valid for the company's current situation, which would be the basis for the rest of the project and allow to achieve the third goal, test alternative scenarios, and evaluate their results.

These three objectives were translated into the schedule in figure 1.

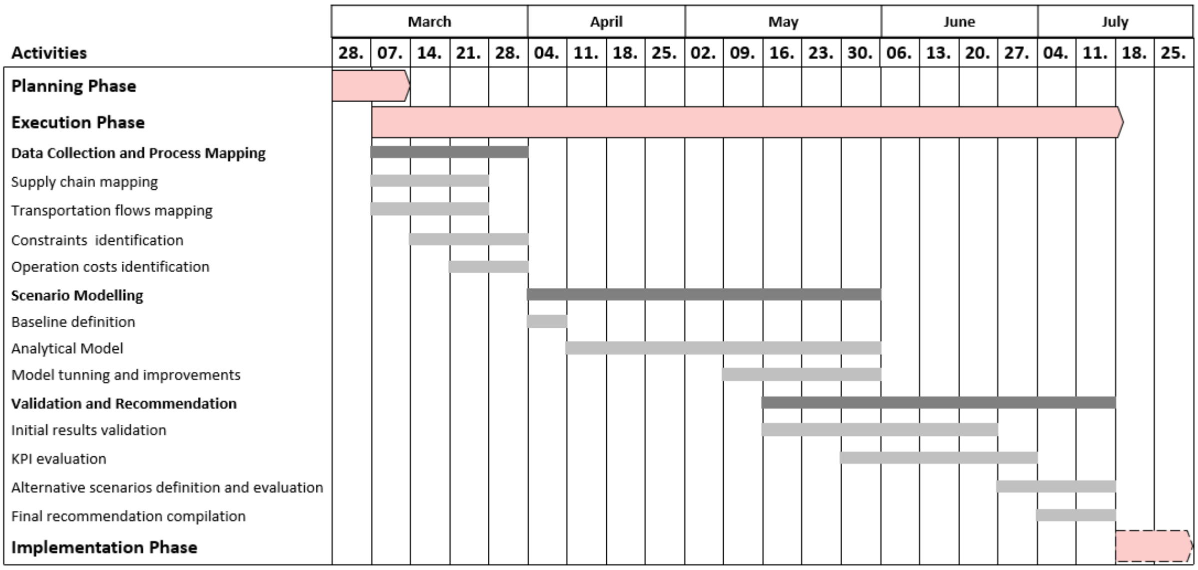

Figure 1 - Chronogram of the project

To better achieve these objectives and guarantee a total alignment between the team and all the stakeholders affected by the possible changes, it was created an iterative methodology that oscillated between two stages: 1) definition of a large subset of routes that already fulfilled a group of constraints; and 2) optimization of new time windows based on the routes given to the model, and the rest of constraints not incorporated in the first phase. After this last stage, there will be a validation phase with the stakeholders to understand what can be done to improve the results and tune the subset of routes of phase 1), moment in each the iterative approach starts again.

It is important to note that even if the proposed methodology was developed and implemented by a team, the author of this thesis was present and active during all phases of the project, with a special incidence in the process mapping, route generation, optimization model conceptualization and validation phases.

1.4 Structure of the dissertation

The current thesis is organized into six chapters to better understand the problem and methodology used. An in-depth theoretical analysis of the vehicle routing problem and the current state of the art of time windows assignment optimization can be found in Chapter 2, focusing on specific research relevant to the method presented in this thesis, while Chapter 3 presents a description of the real-world problem used to validate the proposed model. Next, Chapter 4, describes the mathematical formulation of the optimization model used and the proposed solution approach. Its results and analysis are described in Chapter 5. At last, in Chapter 6 the main conclusions of this thesis and a list of next steps and future work can be found.

2 State of the art of Vehicle Routing and Time Windows Assignment Problems

The present thesis studies an alternative formulation for the Time Window Assignment Vehicle Routing Problem, also known as TWAVRP, that, as the name indicates, is already a variant of the more general Vehicle Routing Problem (VRP) firstly formulated in Dantzig and Ramser (1959). This has since become one of the most studied optimization problems by the research community, especially in recent years when globalization has created new tensions in the supply chains worldwide.

2.1 Vehicle Routing Problem

The original formulation of VRP was called The Truck Dispatching Problem and was considered by the authors as a generalization of the Travelling-Salesman Problem (TSP). Dantzig and Ramser (1959) intended to minimize the mileage of a fleet of gasoline delivery trucks by assigning stations to specific vehicles so that all demands were fulfilled. While at the time, the model couldn’t be tested in a real-world application, the experiences made proved that it could obtain a result (called “best optimum”) very close to the “true optimum” and it was believed that the difference between this two values would decrease with the incrementation of the number of nodes to be connected by the delivery network.

Since then, new challenges demanded new formulations for these problems. Of all of them, three main ones have been the focus of the research community: periodic deliveries, consistency between them and the existence of time windows.

It is also important to mention that all of them might be of interest for the research at hand: stores need perishable goods to be delivered with a certain periodicity, there might exist the necessity to guarantee consistency in those deliveries and as it has been mentioned the assignment of time windows is the main goal of this thesis.

However, the most important formulation for this thesis is the TWAVRP, in which time windows are not constraints, as it happens in the Vehicle Routing Problem with Time Windows, VRPTW, but are a decision variable that we want to optimize.

2.1.1 Periodic Vehicle Routing Problem

Christofides and Beasley (1984) were the first to study this kind of problem. Their periodicity formulation was based on an initial choice of delivery days for each customer to assure the compliance with the service level, and later an interchangeable heuristic that looked for a reduction in the distribution costs that, according to the authors, could achieve a cost reduction of 13% when compared with the previous best-known solutions.

From the few exact approaches of PVRP, Baldacci et al. (2011) created a three-phase formulation: a linear programming (LP) relaxation to obtain a near-optimal solution, a generation of all optimal solutions via an integer problem and, finally, a Mixed Integer

Program (MIP) solver to get the final solution. It was important since it was the first to obtain the best-known solutions for instances used since 1984.

However, periodicity is viewed in the literature from two other points of view. Cordeau et al. (2001) formulated it by defining the space/time between two deliveries to each customer, while Gaudioso and Paletta (1992) limit that spacing using a lower and an upper bound.

2.1.2 Consistent Vehicle Routing Problem

Another problem that logistical operators face in real life is customer satisfaction, especially regarding consistency between deliveries. When a retailer visits a certain client multiple times, they expect to be visited at the same time or by the same driver, which brings new constraints to the VRP problem.

For that reason, Groër et al. (2009) formulated the Consistent Vehicle Routing Problem, also known as ConVRP, as a MIP, by improving a previous algorithm used by Li et al. (2005) to handle very large instances of VRPs. According to him, consistency would be obtained by minimizing the number of different drivers that visited a certain customer and reducing the maximum gap between the earliest arrival and latest delivery to each client while reducing the transportation costs and complying with all other classic constraints of VRP problems.

Later, Ridder (2014) proved that some of the solutions of the original research were infeasible and provided the necessary corrections by including the service times into the total route duration. Furthermore, the author included the ability of a driver to have waiting times during a route, decreasing the maximum difference between arrivals, while having a small impact on the total driving time.

Tarantilis et al. (2012) improved the results for all benchmark instances known at the date by implementing a Tabu Search (TS) to improve the template routes from which the daily ones can be chosen. This generated solutions with smaller total driving times and using fewer vehicles, in a very competitive computational time.

A new variation of ConVRP was introduced by Kovacs et al. (2015) in which instead of creating template routes to be used in the master plan, the authors use a Large Neighbourhood Search (LNS) to generate routes, by removing customers from existing routes and inserting them into better positions. This improved the results from the instances used in previous studies.

Subramanyam & Gounaris (2016) resolve the consistent TSP problem via a branch-and-cut exact solution approach. Their formulation chooses routes with minimum costs for each vehicle and assigns what customers are to be visited. Then, the authors present three different MIP formulations and introduce new inequalities for the problem. One year later, the same authors studied the implementation of idle time during the routes, as Ridder (2014) had done, and the impact it would have on the service costs. According to Subramanyam & Gounaris (2017), the ability to have idle times would cover the increase of 40% in the operational costs of ensuring consistency in deliveries.

Stavropoulou et al. (2019) studied the ConVRP with profits. For this, beyond the frequent customers that the network must serve without exceptions, the authors include a set of potential customers that carry with them a certain profit to the retailer. The proposed formulation was validated with the benchmark instances of Groër et al. (2009) and later extended to the potential clients. The results show that for almost all tested instances, the model can serve an optimal set of potential clients while ensuring the consistency of regular ones. However, the authors also highlight the existence of a “trade-off between the net acquired profits and the service consistency constraints", so all results should be carefully analysed.

Xu & Cai (2018) tried to solve the ConVRP with a variable neighbourhood search algorithm instead of the template route formulation of Groër et al. (2009). This has become the state-of-the-art formulation for ConVRP, at the time, since the solutions obtained (in total travel time and computational time) are better than other papers like Groër et al. (2009), Tarantilis et al. (2012) or Kovacs et al. (2015).

2.1.3 Vehicle Routing Problem with Time Windows

Of these three versions of VRP, the Vehicle Routing Problem with Time Windows (VRPTW) is not only the one with more research published but also the most important for the problem that we intend to study since it was the first VRP formulation to consider time windows.

In Solomon (1987), we can find research into the advances in VRPTW at the time and a compendium of useful heuristics to solve this problem. To analyse those heuristics, the author developed some different classes of instances that are still used as a benchmark to compare different formulations. However, since the research in this field is vast, we shall focus on more recent advances in these formulations.

Taillard et al. (1997) developed a formulation that not only looks at minimizing the travelling cost of a fleet but also includes a penalty for deliveries outside the time windows of each customer. While originally, the model was formulated to only include soft time windows, with small penalties, hard time windows can be implemented by incrementing the cost incurred. The proposed method starts with a random insertion heuristic to generate initial solutions that will later go through a TS heuristic before being stored in the adaptive memory, from which new solutions will be generated. Also, a subset of routes will be combined to generate new routes. All this process ends when the stop criteria are met, and a post-optimization method is applied to improve the solutions.

A few years later, Cordeau et al. (2001) proposed a similar method to the previous one that only varies in the post-optimization phase in which it uses a special heuristic developed for the Travel Salesman Problem with Time Windows. While the results are not better than the previous ones, this methodology has an advantage in terms of flexibility, memory usage and computational speed.

Pisinger & Ropke, (2007) developed a methodology that can be adapted to five different types of VRP problems by transforming all problems into "a rich pickup and delivery model" and solving it with an Adaptive Large Neighbourhood Search (ALNS) framework. When applied to the Solomon (1987) instances, it presents better results than 4 out of 6 previous models, while the two remaining don't report the number of experiments needed to obtain such results. When compared to the exact methods, this heuristic could get very close results, if not find the exact optimum solution.

Kallehauge (2008) reviews four different formulations of VRPTW and the exact associated methods of each one. While the path formulation is the most common and successful of the four, the spanning tree formulation still needs to be proven. Of the other two, the arc and arc-node formulation, the arc formulation is considered the more promising one and where further study can achieve better results. Four years later, Baldacci et al. (2011) produced another review of exact algorithms for VRPTW.

In Taş et al. (2013), a new formulation is made in which the objective function includes a larger set of costs: not only the travelled distance, but also the dimension of the fleet, the driver's overtime, and penalties for early and late arrivals. These authors also propose a TS method and consider the stochasticity of travel time. While the results, when compared with other proposed methodologies, are good, the most important contribution of this research is a managerial insight: "with a small increase in the total transportation cost, we observe a significant decrease in the total service cost". This study was improved by Taş et al. (2014a) with an implementation of an ALNS heuristic that benefits the problem when it considers the service time of the delivery. Later that year, the same author introduced a new type of VRPTW, with flexible time windows and more restricted solution space in Taş et al. (2014b), in which a vehicle has a tolerance to arrive outside the delivery window of a certain store. This flexibility brings operational cost savings for the majority of the Solomon instances since it allows for a reduction in both total distance travelled and the number of vehicles used.

While the preliminary results significantly increased fleet size for the benchmark instances, it is important to mention a formulation for a VRPTW with stochastic travel and service times with a different approach: the use of a multi-population memetic algorithm by Gutierrez et al. (2016).

A more common strategy to solve VRP problems is a branch-price-and-cut algorithm used by Errico et al. (2016), in its formulation of a VRP with hard time windows and stochastic service times. Even using two different recourse policies, this formulation works better in instances with smaller service time uncertainty and for up to 50 customers.

2.2 Time Window Assignment Vehicle Routing Problem

Of all VRP formulations known, the Time Window Assignment Vehicle Routing Problem, known as TWAVRP, is the one that allows the allocation of time windows to reduce transportation costs. According to Neves-Moreira et al. (2018), TWAVRP is closely related to a more known optimization problem called VRPTW and presents "at least, the same level of complexity". This problem was firstly formulated by Spliet & Gabor (2015) and since then there were only written twelve articles about it, according to Jalilvand et al. (2021).

In 2015, Spliet & Gabor (2015) started studying a new optimization problem to be applied in cases where "time windows have to be assigned before demand is known. Next, a realization of demand is revealed, and a vehicle routing schedule is made that satisfies the assigned time windows", that they called TWAVRP. This problem is very common in retail operations in which a retailer with a large set of stores needs to define the timing of delivery of a product to each store before knowing the quantities that the store will order. Normally, this time slot is agreed upon between the two stakeholders and must be complied with daily, independently of all surrounding factors. For this, the authors defined two different kinds of windows. While the exogenous define the period $[a_s, b_s]$ where the delivery can be made, the endogenous is the exact delivery window. For example, suppose a delivery can only be made during the opening hours of the store (exogenous). In that case, there are infinite endogenous time windows inside this exogenous period that can be chosen.

The model was based on the ConVRP formulation by Groër et al. (2009), with some differences. In TWAVRP all customers require the delivery of goods every day, both supplier and customers belong to the same logistical network, and, finally, due to the high volumes of goods being transported for each customer, the capacity of the vehicle will, in essence, prohibit the exceedance of maximum driving hours.

This initial formulation of TWAVRP consists of a column generation heuristic to calculate lower bounds and a pricing problem to be solved via the algorithm proposed by Ioachim et al. (1998). Finally, the optimal integer solutions are obtained using a branch-price-and-cut algorithm.

The first phase or the column generation algorithm intends to resolve a linear programming relaxation of the formulation. At each iteration, only a subset of routes is used, followed by the pricing subproblem that intends to find feasible routes with negative reduced costs that can be added to the subset of routes being used. When no routes can be added, the current solution is considered optimal for the first phase of the process. Then, the authors propose a branch-price-and-cut algorithm that obtains the upper bounds for the problem when a solution with integer arc flow in each scenario is found to the LP relaxation.

This first formulation worked for up to twenty-five customers and allowed for a cost reduction of up to 5.42% when compared to a VRPTW with average demand. However, the authors defined that further study into this area would be favourable to allow larger instances and the use of better inequalities.

In that same year, Spliet & Desaulniers (2015) proposed a discrete version of the same problem – Discrete Time Window Assignment Vehicle Routing Problem or DTWAVRP –, considered more applicable for real-world problems, that when applied to five more customers and two more demand scenarios than the original allowed for a cost reduction of 3.64%. Instead of allowing a time window to be selected from a continuous interval within the exogenous time window, the discrete formulation considered "for each customer a finite set of candidate time windows from which one has to be selected". Another novelty that the authors brought with this paper implementing a Tabu Search (TS) to accelerate the column generation algorithm, especially to solve the pricing problem of creating new routes.

These two formulations assume that the travel times between two nodes are constant, which according to Spliet et al. (2018), is a "very short-coming in many real-life applications". With this model, the authors intend to include in the selection of time windows the predictable variations in travel time, such as the congestion of rush hours in urban areas. To do so, and based on the previous work of Ichoua et al. (2003), they developed a piecewise linear continuous function $\tau_{ij}(t)$ that represents the travel time between two nodes *i* and *j*, when *i* is departed at time *t* and included it into the extended mathematical formulation of the previous algorithms. In the same year, Dalmeijer and Spliet (2018) also developed the best-known procedure for instances with only one product, with an exact algorithm that reduced computational time and allowed for the inclusion of more customers than all previous studies. This state-of-the-art algorithm is 6.6 faster than the original TWAVRP. However, it doesn't guarantee a consistent optimal solution for instances with more than 35 customers.

In that same year, a new algorithm was formulated "to solve the sampled deterministic equivalent" of the TWAVRP stochastic model by Subramanyan et al. (2018), with a decomposition of scenarios. The authors of this study defend that decomposing the TWAVRP formulation onto its scenarios allows solving the problem only by tracking the applicable time windows to each node of a branch-and-bound search tree without an explicit encoding of the assignment of time windows. When compared with the benchmark of Dalmeijer and Spliet (2018) and Spliet and Desaulniers (2015), this deterministic version of the model allowed to solve more unsolved instances by the previous formulations with the same or faster computation times.

Martins et al. (2019) studied the impact of incorporating multi-compartment vehicles into the retailers' fleets. With these vehicles, two or more different product segments can be consolidated into the same delivery, which did not happen in most of the literature around TWAVRP that only considers one product segment. This will have a huge impact on the logistical operations (increasing the complexity of the loading operations in the warehouse and decreasing the number of deliveries to a store) but adds a new tactical decision of which product combinations should be transported together. The proposed Product-oriented Time Window Assignment for Multi-Compartment Vehicle Routing Problem, called PTWA-MCVRP compares well with other benchmark datasets. It presents gaps close to zero and solutions that respect the consistency clauses with a run time acceptable for tactical planning problems.

Li et al. (2019) adapted TWAVRP with both constant and time-dependent travel times to a bi-objective version with the intention of not only minimizing travelling costs but also maximizing the customers' satisfaction level, known as the bi-objective time window assignment vehicle routing problem (BOTWAVRP) and bi-objective time window assignment vehicle routing problem with time-dependent travel times (BOTWATDVRP). For the model to work, each customer will have at their disposal a set of discrete non-overlapping time windows that cover an entire day, from which they can choose two as the preferred ones. Implementing a hybrid multi-objective genetic algorithm has a positive effect on the computational results being able to solve all instances, always under two minutes, and finding a "satisfactory near-optimal front at a single run".

Wang et al. (2020) formulated a problem which considers a collaborating multi-depot logistics network allowing a customer to be served by another distribution centre that is not its own. The study proves that the reorganization of customers per centre and the cooperation between different depots will allow for significant cost reductions.

Contrary to all previous research, M. Jalilvand et al. (2021) proposes a two-layer time window assignment vehicle routing problem with stochastic service times (2L-TWAVRPSST) in which an endogenous time window has two layers. The inner layer is nice to have, while the outer layer works as a hard constraint. Servicing inside the inner layer is the preferable solution while servicing outside the inner layer but within the bounds of the outer layer is also acceptable with a corresponding penalty. Then all costs are included in the objective function, and a method based on Progressive Hedging is applied.

2.2.1 Time Window Assignment Vehicle Routing Problem with Product Dependent Variables

Neves Moreira et al. (2018) formulated a version of TWAVRP for the selection of different time windows for the same customer, according to distinct product segments, called Time Window Assignment Vehicle Routing Problem with Product Dependent Variables or TWAVRPP.

According to the authors, the food retail sector is especially affected by the existence of this segment, since its goods are palettized at different times according to the time they reach the depot or the temperature they need to be transported at. While fruits and vegetables must be transported at a controlled but positive temperature, frozen food requires conditions below zero degrees, and canned food or cereals can be moved at environmental temperature. Moreover, the way these products can be stocked inside the stores implies that all refrigerated food must be delivered before the store opens. In contrast, other food can be delivered during the day since customers expect to find fresh food early in the morning. With all these constraints in mind, if a model didn't allow for a product segment dependency, all deliveries would be made in the early hours of the day, implying the necessity of a large fleet to operate during a short period of the day.

Since the author intended to develop a method that would allow for a real-world application and tackle instances with hundreds of stores to serve, a hybrid method divided into three phases was chosen.

Firstly, a route generation is performed to create a large set of routes to be used in the subsequent phases. Three kinds of routes were considered: direct routes (1 product segment is delivered to 1 store), multi-location routes (1 segment to several stores) and multi-product routes (multiple segments consolidated to only one store). For this, three approaches are recommended: the nearest neighbour heuristic via cheapest insertion, an Adaptive Large Neighbourhood Search (ALNS) procedure, and a historical set of performed routes used by the retailer to validate the method. All routes generated by these methods were later replicated throughout the exogenous time window.

Secondly, to schedule the deliveries, an initial solution is constructed through a series of sub-problems and the use of a Relax-and-Fix heuristic first proposed by Pochet and Wesley (2006). This heuristic is based on portioning the routing variable into three sets (fixed, optimized, and relaxed). In the beginning, all routing variables are relaxed, and a certain group at the time are optimized. After a value is obtained, they are fixed for the next iterations, where another group is optimized and later fixed, until all variables are finally fixed.

At last, the improvement heuristic used is a Fix-and-Optimized (FO) first proposed by Helber and Sahling (2010). To fight against the computational effort needed to compute the number of integer variables, smaller subproblems are created and solved by a Mixed-Integer Programming (MIP) solver.

The biggest advantage of this approach is the consolidation of different product segments into one delivery. For that reason, when applying this method to instances provided by a large European food retailer, the existence of multi-product routes allows for an average travelling cost improvement of 6.44% when compared to the exclusive use of single product routes. The comparison with the formulation of TWAVRP by Spliet and Gabor (2015) allows an improvement of 5.3%.

While the results seem very favourable, and the methodology allows for a closer approximation to the reality lived by food retailers all around the world, two big problems can be identified and these are the starting point of the current thesis: 1) a necessity to add more constraints related to the warehouse loading operations and in-store workflow management, and 2) the general overlook of other logistical costs that these operations are subject to besides the travelling cost between two nodes.

3 Description of the case study

The current project was developed in a national food retailer that manages a daily delivery operation of edible goods between two main warehouses and more than three hundred stores. The project's main objective was to define a new set of delivery time windows for each of the six product segments and to each store, having in mind two important dimensions: logistical costs and entropy made by a change.

Changing the time window in which a group of products is delivered to a store has not only an impact on the transportation costs (as most of the literature defends) but also a massive impact on the logistic operations of the warehouses and the day-to-day workflow in the store. For that reason, only the implementation of a TWAVRPP that considers the minimization of all logistic costs, and not only the transportation parcel, but also all operational constraints will generate feasible solutions that all stakeholders of the process accept.

A graphical representation of the supply chain of a food retailer can be found in figure 2, where the main operational tasks of the three main stakeholders are identified. These three stakeholders are the warehouse team, the planning team responsible for the transportation plan and the store managers.

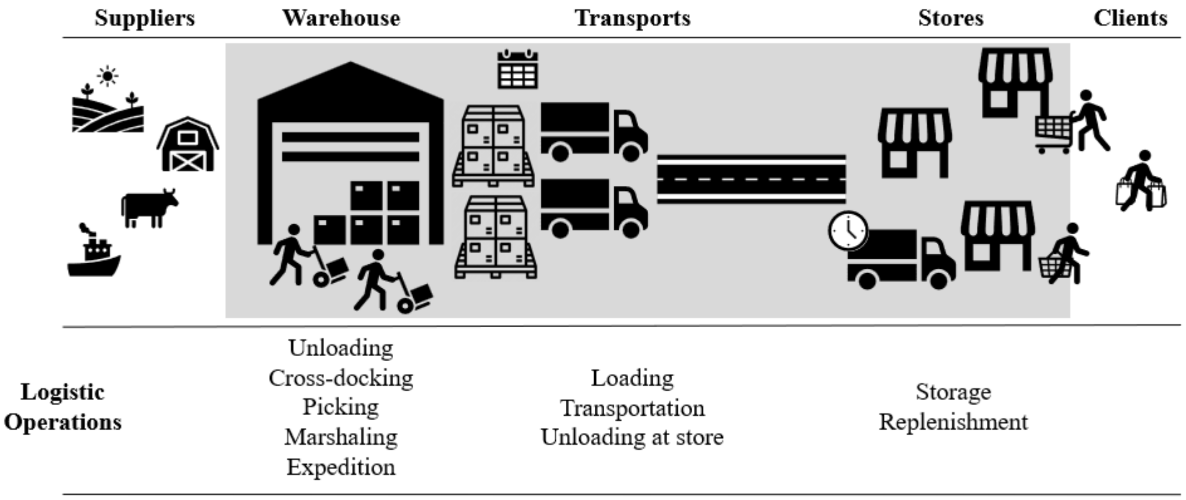

Figure 2 - Main operational tasks of the three main internal stakeholders

The daily operation of replenishing goods to stores has 5 main stakeholders, 2 external and 3 internal. From the externals, the suppliers have a smaller sway in the operations and are only responsible for delivering their products to the retailer's warehouses. The clients, the second external stakeholder, are the main source of revenue for the retailer and, for that reason, the ones whose opinions must be more safeguarded. This means that any solution provided to the retailer has to guarantee the well-being of the final customer and that they can always find the goods they are looking for at the store.

The other three stakeholders are internal to the operation and the ones who facilitate that the product arrives from the supplier to the client. While they must adapt to the best solution for the customer, they are the ones who present more constraints because their resources are finite.

3.1 Warehouse

The operation starts at the warehouse, where two tasks can be performed after receiving the products from the suppliers. If the products need to be sorted into pallets, then a picking operation occurs. However, if the pallets arrive already arranged by final location, then the depot workers only need to process a cross-docking. When a pallet is finalized, it can be transported to the marshalling area until it can be expedited and loaded into the trucks for transportation.

However, the main trait that impacts this operation is the kind of segment to which the product belongs. To facilitate the analysis, six main segments were assumed based on three factors: temperature, picking operation, and the product origin (table 1).

Table 1 - Mapping of product flows

Temperature	Picking	Origin	Segment
Ambient	Pick-by-Line	Directly from supplier	**PBL**
	Pick-by-Store	Directly from supplier	**PBS**
Controlled or refrigerated	Pick-by-Line	Directly from supplier	**Fresh goods**
	Cross-Docking	Smaller Warehouse inside the network	**Meat**
	Cross-Docking	Smaller Warehouse inside the network	**Fish and Cod**
Negative	Pick-by-Store	Smaller Warehouse inside the network	**Frozen goods**

From this division, the company decides to which part of the warehouse the supplier, or the previous node of the network, must deliver each product. Everything varies from there on.

According to this separation, the following variables are strictly dependent on the product segment, even inside the same warehouse compound:

- The time of day in which the products arrive from the supplier.
- The amount of time to build and transport a pallet.
- The cut-off time (the most important variable) in which the pallets are ready to be marshalled and loaded.
- The capacity of the marshalling area in pallets.
- The number of docks available for the loading operations.
- The resources available, mainly labour force and picking and loading equipment.

For example, PBS and frozen goods are ready to be shipped throughout the day (due to the PBS picking operation). In contrast, PBL and fresh foods directly depend on the arrival of the last supplier, while fish and meat depend on the arrival schedule from the respective warehouse facilities. Furthermore, it is also not the same to build pallets of canned goods or square boxes of cereals or do it with fresh food which needs to be handled with much more care.

Moreover, the fresh food flow, more perishable than the rest, is more constrained by the arrival time at the store and the necessity to ensure that it is always in good condition and ready when the stores open, so that the first clients of the day have already the products at their disposal.

For the warehouse manager, the optimal time windows would allow him to reduce the workforce and load trucks during the day when the labour is cheaper. It would probably be better to fill trucks with only one product flow to facilitate the operations and minimize loading time and the amount of time a pallet must wait in the marshalling area.

3.2 Planning and Transportation

The planning team is responsible for receiving the stores' orders from their internal information system. All individual SKUs that were ordered will be palletized according to their product segment and store to be delivered, which will be called a service.

From this point onwards, the planning team only focus on the number of pallets of each segment that each store has ordered and translates them into a schedule of truckloads that the warehouse team must prepare.

This schedule must ensure that all cut-off times are ensured and that each truck arrives at each store respecting the specific delivery time window for each flow. Another necessity of this schedule is to assign a specific truck with a semi-trailer or a rigid truck according to the capacity needed and the conditions to access each store on the route. Finally, the schedule must obey the consolidation rules of each warehouse that govern which flows can go together inside what vehicles, as we can see in table 2.

Table 2 - Mapping of kind of trucks necessary for the operation

Consolidation of flows according to temperature			**The minimum type of semi-trailer or rigid truck required**
Ambient	**Controlled**	**Negative**	
✓			**Isothermal**
	✓		**Mono temperature**
		✓	
✓	✓		
✓		✓	
	✓	✓	**Bi temperature**
✓	✓	✓	

The ambient products (that were previously divided into PBL and PBS) can be carried in all kinds of trailers since they don't need a specific temperature to be transported in. This also means that they can be consolidated with all other product segments.

The segments that need a controlled temperature (fresh food, fish, and meat) can only be consolidated with frozen food if the trailer has two temperature engines. If not, they must be separated into mono-temperature trailers or rigid trucks.

Knowing these rules, the planning team must prepare a plan by associating a truck with one or more stores and one or more flows to guarantee that all existing demand is delivered at the smallest cost possible. With this plan, the planning team can generate four distinct types of routes that the company vehicles can do, as mapped in table 3.

However, this planning is not as simple as it might seem. Not all stores can be visited by all kinds of trucks, and some stores have specific constraints that restrict the combinations that can be made. Moreover, the fleet that a planner has at their disposal is finite with a certain number of trucks and drivers available.

Table 3 - Mapping of kinds of routes used in planning

Number of stores	Number of flows	Type of route
1	1	**Direct route to the store**
1	> 1	**Consolidated route**
> 1	1	**Circuit**
> 1	> 1	**Mixed route**

By adding all of this to other specifications such as the dimension of the fleet (number of trucks, rigids, and semi-trailers), the time it takes to load them, the ability to sign external occasional help from logistic companies and special contracts to do specific journeys, then the initial flexibility of daily planning will decrease substantially.

When considering the number of constraints that the transportation phase of the logistical cycle is subject to, it is easy to understand why most authors mentioned in chapter 2 only consider these costs in their objective function. However, there are still more rules that a planner must have in mind. The labour code and employment regulations govern the number of hours a driver can drive during a day, the amount and size of pauses, and other aspects that conflict with the maximization of the fleet utilization rate and consequently increase the transportation costs.

For the transportation manager, the optimal time windows would allow a reduction of the fleet costs, with a constant number of vehicles travelling across the day and a utilization rate closer to 100%. Other preoccupations would be reducing the total kilometres driven, the total number of travels, and the number of vehicles subcontracted to logistic partners. If possible, the transportation manager would prefer to change windows according to the demand of that day and the possibility to anticipate services to minimize its costs.

3.2.1 Suppliers

To optimize the space of the vehicles when returning to the warehouse from the last delivery of the route, the retailer developed an interesting strategy that also needs to be accounted for. That strategy consists of selling that space to some suppliers so that the goods are transported to the warehouse by the retailer’s fleet. This has advantages for both the supplier and the retailer. The retailer optimizes the useless travel time of the vehicles and profits by selling this service to the suppliers without incrementing significantly the total travelled distance and the operation costs. The supplier, instead of having its fleet, pays for the retailers’ fleet, incurring a smaller cost.

This reality needs to be accounted for by guarantying that the routes performed include the incremental distance and travel cost associated with visiting the supplier after delivering the last pallets to the last store.

3.3 Stores

The final stakeholder involved in all this operation is the stores. The company owns four main insignias of stores that differ according to size, location, and concept, as is depicted in table 4.

However, most constraints depend directly on each store and not on the concept it belongs to, which means that each of the more than three hundred stores has a unique set of characteristics and physical properties that allow for deliveries at separate times, by different vehicles and under different conditions.

Table 4 - Mapping of retailers' store concepts

Kind of store	Description
Hypermarket	Urban premium locations at shopping centres
Large Supermarket	Large dimensions and a vast range of products
Supermarket	Urban locations to answer to everyday needs
Minimarket	Franchised proximity stores in urban and rural areas

The delivery time windows directly affect a store and the replenishment of goods on shelves. The replenishment of goods is made by the workers of each store and implies that, when they are available to do so, the product is already unloaded from the truck and, if needed, be stored somewhere inside the store. If the quality of the product depends on the temperature they are stored, it is necessary to have free space in refrigerated chambers to accommodate them. Another important thing to have in mind is the physical access to the store, which many times does not allow any type of vehicle to deliver the products. The type of pier and loading material available in each store can also condition this decision, as are the alarming systems.

Finally, there are some of the external constraints that a store is subject to. Its location inside densely populated urban areas might impose time limitations due to noises during the night, dangerous neighbourhoods with high crime rates, or lack of infrastructures or personnel to help. Even the existence of hospitals, schools, or construction work nearby can create huge implications for the definition of time windows and the fulfilment of time of the entire plan for the day.

To fight some of these factors, the retailer has implemented in some of its stores the concept of autonomous deliveries, which allows delivery during the night to some stores with certain dock conditions and the existence of a refrigerated chamber to store the perishable goods. This type of delivery also implies the necessary equipment, the ability to access the chambers and a helper to accompany the driver if the store doesn't allocate a worker to receive the goods at the delivery window.

Normally, a store manager would choose delivery windows that would allow the fresh food to be sorted just before the store opening time, and all other flows to be delivered throughout the day to allow the workers to replenish the shelves during their normal working hours and never have stockouts of products or lack of workers available in the store. Normally, a store manager would not like a window to vary according to the day or the demand for that day.

3.3.1 Collecting at the store

The real-world logistic operations of a food retailer include not only daily deliveries of products to stores but also an increasing necessity to collect empty pallets, tares, recipients and so on. This necessity is created by growing sustainability concerns, but also less available space at the stores to store such materials.

However, a vehicle can only load such goods at the last stop of the route, to not block access to the goods being transported. For that reason, the performed routes selected and corresponding optimized time windows, need to ensure that the stores that nowadays have the necessity of collecting tares continue to have such needs met.

3.4 Cost drivers and KPIs

Another important analysis that arose from the supply chain flows was the identification of all relevant cost drivers behind the logistic operation in all three stakeholders and the main KPIs that each one valued so that the final recommendation would be acceptable to all of them. It is important to mention that while some cost drivers and KPIs are strictly related, some are not, as is depicted in figure 3.

When looking at the warehouse, the main cost was the labour associated with the planning, so it was important to calculate the loading FTEs required to execute a new plan with the proposed new delivery time windows. It was also important to control the number of pallets being loaded into trucks at peak times, the number of loadings of the fresh flow at peak since there is an identified bottleneck in this flow and the number of total loadings.

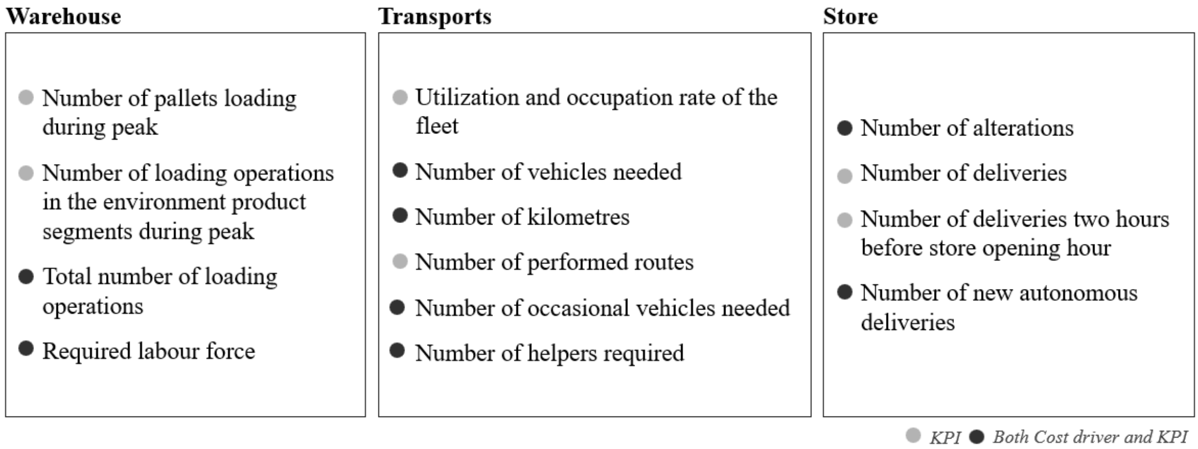

Figure 3 - Main cost drivers and KPIs to be considered when evaluating new sets of time windows

Identifying the costs behind the transportation phase was easier because that is the area with more research done when analyzing similar VRP problems. It was important to measure the number of kilometres travelled, number of vehicles needed, number of subcontracted trucks, and the number of autonomous deliveries made in the new time windows. Together with the team behind the planning and transportation, it was necessary to identify the variable cost per kilometre, fixed cost per vehicle, costs of an occasional external service, cost of a helper for the autonomous deliveries and to obtain a distance and duration matrix between all relevant nodes of the operation. To evaluate the quality of the results other metrics were necessary such as the fleet utilization and occupation rates, and the total number of travels.

When looking at the impact of changing the delivery windows on the daily operations of the stores, the monetary impact is only relevant if there is a need to equip a store with equipment or facilities that it does not have at present. There might also be a necessity to adjust the labour schedules, but in the long run, that impact might be neglected.

It was also considered a cost of changing the delivery windows to reduce the ability to change a window with small financial benefits to the entire operation since negotiating a new window with each store might be very time and effort consuming. The store's managers also want to control the number of deliveries they receive, especially, the number of deliveries done 2 hours before the store opens.

All these costs were considered OPEX (Operational Expenditures) except for the investment necessary to give the conditions to allow any store to receive autonomous deliveries during the night, which is CAPEX (Capital Expenditures).

3.5 Solution Method

Such a complex supply chain originates two important questions to tackle. The number of stores will generate instances that might be too big to be solved by some of the exact methods reviewed in Chapter 2, that as it was told, could only solve instances up to 40 customers. The amount and complexity of constraints might also bring computational challenges.

For that reason, a hybrid approach might be the best option to solve this problem and obtain feasible results with an acceptable computational effort.

The proposed strategy includes three distinct phases. A data processing phase that intends to create a finite set of routes to be used to generate an initial solution. In this phase, some of the most important constraints will be included so that all routes that come out of this part are feasible. This intends to reduce the size of the instance and the computational effort of the next phases. The second phase will be an optimization model to obtain the best time windows for the set of routes that minimizes the operational costs. The third phase is validating the results and improving the model. To guarantee that the results obtained are in line with all stakeholders, phases 2 and 3 should be iteratively repeated until everyone is comfortable with the results.

All this process is depicted in figure 4.

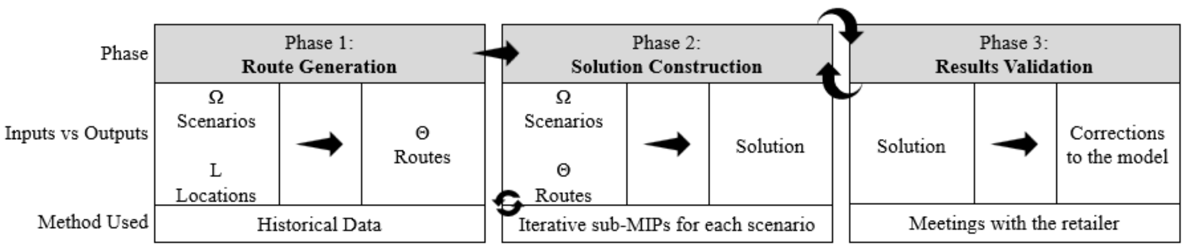

Figure 4 - The three phases proposed

4 Methodology

This chapter will dive into the proposed methodology to solve the TWAVRPP. As was mentioned in section 3.5, to solve large instances it is better to assume a hybrid approach. For this thesis, the best solution was to design a three-phase method that will be further explored next.

Section 4.1 develops the route generation phase which mainly aims to generate the instances to be used by later phases. It also expands on the strategies deployed to reduce the instance size and the computational effort.

Section 4.2 presents the problem statement and mathematical formulation of the optimization model developed for this project so that it can be more easily adapted for another research while also presenting additional constraints that can be easily implemented in future work.

At last, Section 4.3 tries to explain the importance, for this project, of the validation phase and its iterative relation between phases 2 and 3.

4.1 Phase 1: Route Generation

One of the biggest challenges of trying to solve this problem with exact methods is the dimension of the data that would be needed to include in the optimization model. Suppose we consider a set of 100 customers and try to create all possible routes that go directly from the warehouse to those customers, all possible routes that visit two different customers or all possible routes that visit three different customers. If after creating all routes we unfold them throughout the day so that each route can leave the warehouse multiple times during the day so that it repeats every 30 minutes, we would get more than 47 million possible routes, as is depicted in equation 4.1.

$$(Direct\ Routes + Two\ Sotre\ Routes + Three\ store\ Routes) * 24\ hours * 2\ {}^{times}/_{hour} =$$
$$= \big(100 + P(100,2) + P(100,3)\big) * 48 = 47\ 049\ 600\ routes \tag{4.1}$$

If we, then, consider the real case of implementing different product segments, the existence of different vehicle capacities or even circuits with more than 3 stores, then the dimension of this instance would reach unsurmountable values. For that reason, the first phase proposed looks to include all existing operational constraints to the routes so that the set of routes generated and used in the later phases of the approach only includes feasible routes.

The literature reviewed in Chapter 2, provides several heuristics and methods to generate routes such as the nearest neighbour, LNS, ALNS, template-based routes, and TS among others, that could easily be used in this project. Furthermore, there is a clear agreement between most of the authors that the necessary inputs for this phase would be a set of locations of the customers and the corresponding demands.

The close partnership established with one of the largest food retailers in Portugal allowed obtaining a real-world set of demand scenarios and the locations to which the deliveries should be made. Moreover, such a partnership allowed to get the entire set of historical routes performed by the retailer for an entire year. A deep study into this set allowed us to understand that the routes already performed hold a large diversity that would be enough to replace the route generation methods established in the past literature while allowing for disruptive solutions with a guarantee of complying with constraints of the logistical operation.

From this historical set, all four types of routes mapped in table 3 were extracted and each service was served by multiple routes of different types and with different characteristics. To better store all obtained information, a database was designed (as seen in figure 5) to help manage the existing data and facilitate the next steps in this phase. Each route is only performed by one vehicle of the fleet, with a specific capacity and temperature, and visits a sequence of stores totalling a certain distance by a certain cost. Each service will have a certain time window (that we want to optimize) with a certain amplitude, and each scenario will have a certain group of demands associated with it.

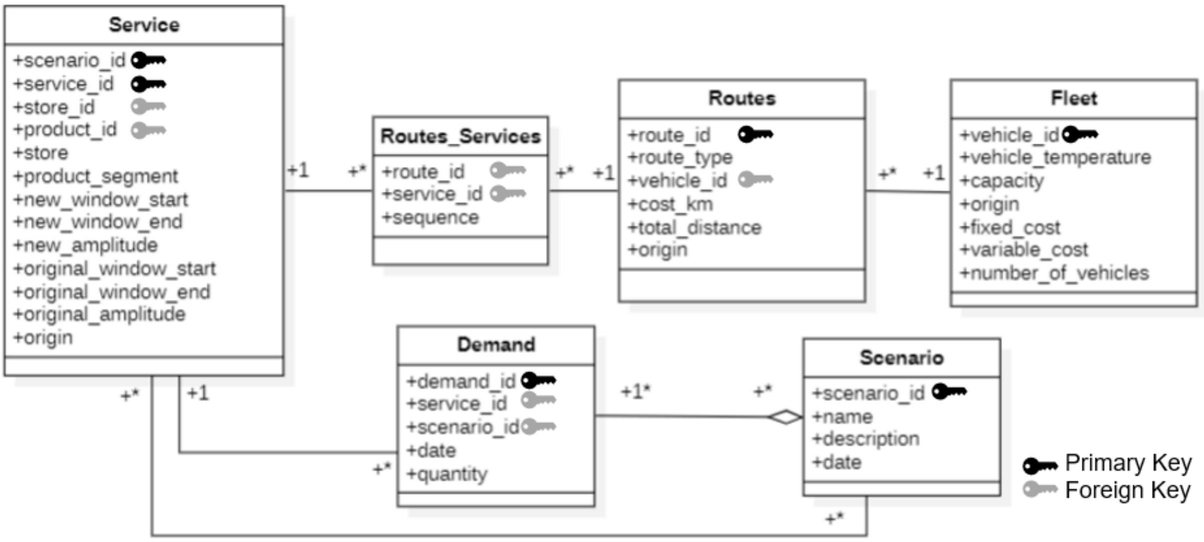

Figure 5 - Initial database structure

While using historical data ensured diversity, the current set of routes only visited each store during the current delivery time windows. If that set was used to optimize the time windows, then we would obtain the same windows that already exist since there are no degrees of freedom to choose different windows. For that reason, after obtaining all routes it was necessary to unfold each route within the planning horizon into 30-minute periods, as represented in figure 6. It was important to collect from the historical set of routes the departure time from the warehouse of each route and multiply it throughout the planning day. Since the time windows' granularity is 30 minutes, this multiplication was sufficient.

To manage this multiplication, it was important to add a new dimension to the database: the time of departure and two new tables that connect such timing with the routes and with the services, respectively.

For each route departing from the warehouse at a certain time, a new id route/hour was created since, according to the time of the day, the same route could incur different costs (as will be further explored in the next section).

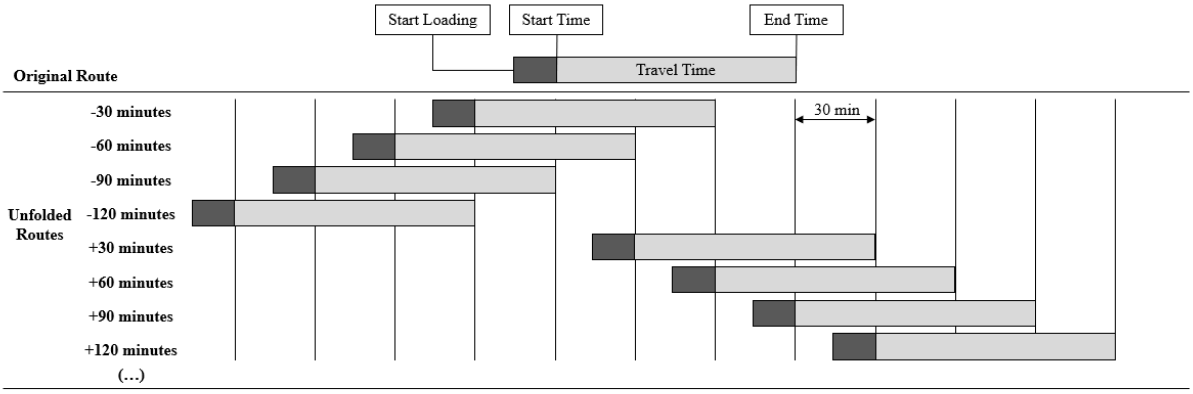

Figure 6 - Visual representation of multiplication of a historical route into 30-minute time slots

For each service, each route departing at a certain time would arrive and leave the store at different times, which need to be preserved, since these are the times that will determine the new optimized time windows, as depicted in the database structure of figure 7.

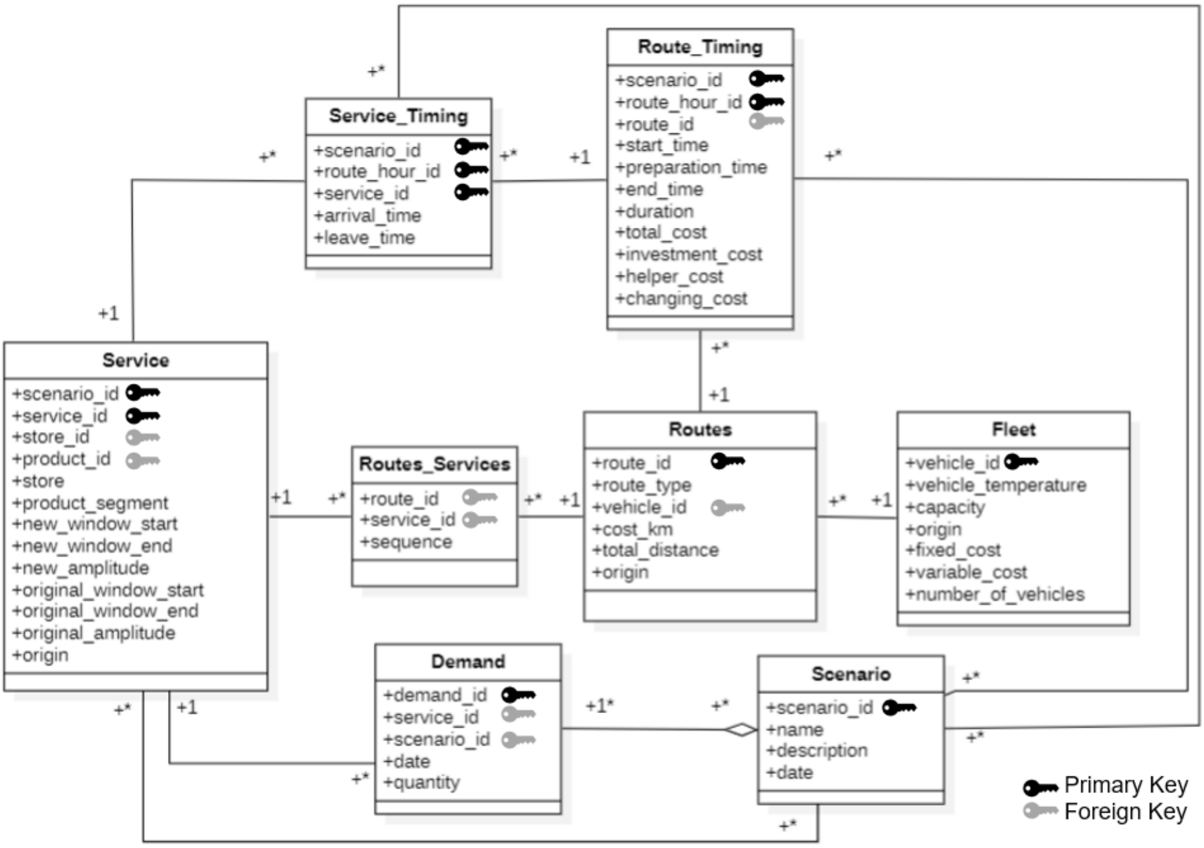

Figure 7 - New database structure

It is important to remember why it was decided that a hybrid approach was the best way to solve large instances of the TWAVRPP: to reduce the number of routes provided to the optimization model by only giving feasible routes.

And here is where one of the biggest advantages of only using historical routes appears. When doing so, we guarantee that all historical routes are already feasible because if they were not, they wouldn't ever be planned and happened. However, some constraints start being ignored when multiplying a past route into the remaining 47 timeslots of departure within one day.

Therefore, it is important to separate the logistical constraints into two different sets. The first one does not depend on the time the route is performed, and so, when using historical routes, is never disobeyed. To this set belong constraints such as the compatibility between the size of a vehicle and a store and all drivers-related matters like route duration and the mandatory pauses.

However, the second set of constraints that depend on the time of the day is the one that needs to be verified in that stage of phase 1 since it is the one that will guarantee the feasibility of all routes that will be delivered to phase 2 of the hybrid approach. In this feasibility check, it is important to include: 1) the cut-off time of each product segment, 2) the arrival time of the fresh category to the store, 3) the existence of hard constraints, that even with investment, don’t allow for an autonomous delivery during the night and, at last, 4) the existence of specific restrictions that impede deliveries during specific hours.

The cut-off time is the time of the day in which a pallet from a specific product segment is ready to be transported and loaded into a vehicle. For each product segment, there cannot be any route in which the preparation time happens before this cut-off. The second constraint limits any routes that deliver fresh products after the opening of any store. The third one intends to limit autonomous deliveries to stores where it wouldn’t be possible, even with an investment. The best example of this constraint is the lack of physical space to install a refrigerated chamber (in case it doesn’t exist), which costs hundreds of thousands of euros and would make these routes unviable. The last intends to identify the stores that can’t receive products at specific times and delete any route that intends to deliver at that time. Regarding the final topic, dangerous neighbourhoods where the crime rates might see the products being stolen, legal injunctions near specific neighbours that complain about the unloading noise during the night, lack of access to docks or the ability to be visited by vehicles with different capacities during different periods of the day are some examples that limit the delivery hours. The proposed methodology can be found in Algorithm 1.

Algorithm 1: Pseudo-code for the route generation phase

procedure ROUTE GENERATION (Ω, L, historical data)

  $\Theta \leftarrow \emptyset$

  $\Theta \leftarrow$ Add routes from historical data

  $\Theta \leftarrow$ Create copies of each route with 30 minutes intervals along exogenous time windows

  for each $r \in \Theta$ **do:**

    $\Theta \leftarrow$ Remove routes with product segments being prepared before the cut-off time

    $\Theta \leftarrow$ Remove routes delivering fresh segment after the opening the of store

    $\Theta \leftarrow$ Remove autonomous routes visiting stores without conditions

    $\Theta \leftarrow$ Remove routes visiting a store at a time in which they can’t

return Θ

An example of this feasibility check for a direct route of fresh food can be found in figure 8.

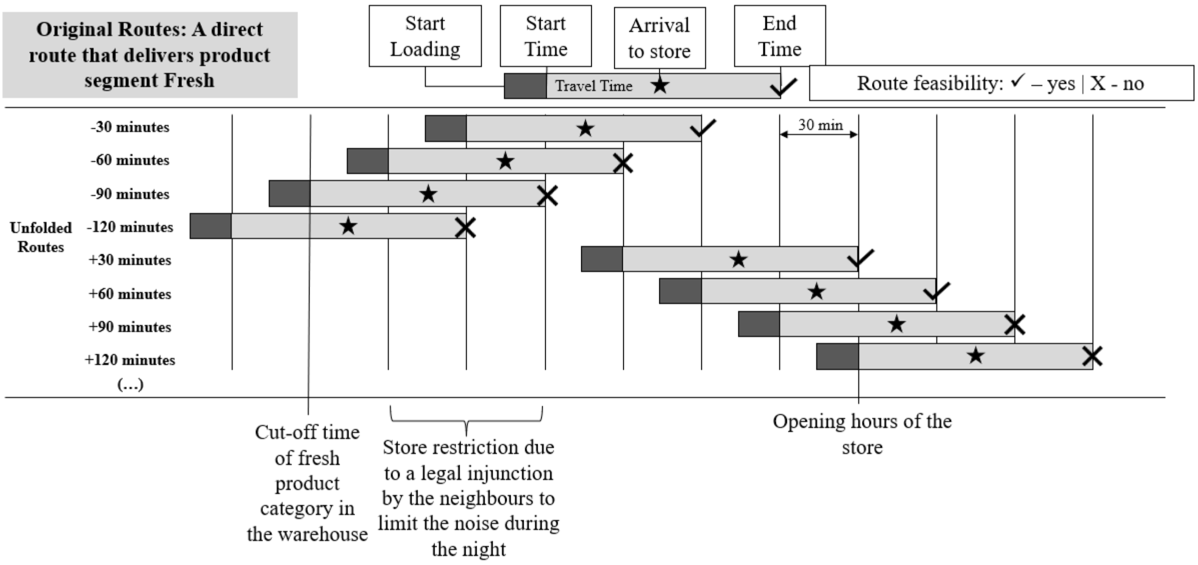

Figure 8 - Example of the feasibility check for a direct route that delivers products of the segment Fresh to a specific store with impact of the store opening hours and the cut-off time

Finally, it was also important to check if in between the execution of the routes and the present day there were assumptions that changed and if so needed to be corrected like stores that ceased to exist, different travel times due to new roads, windows that changed in that period and so on.

4.2 Phase 2: Optimization Model

The challenge of optimizing an existing set of time windows to deliver specific products at different times to each store was initially formulated by Neves-Moreira et al. (2018). However, there are changes to be made to adapt better the theoretical model to the real world. Therefore, the present thesis will adapt said model with changes proposed by the model's creator in that same paper, but also with other alterations that were studied and drafted during the development of this project.

This section presents an optimization problem formulation that receives the set of routes generated in the previous phase and obtains a set of optimized time windows for each scenario.

4.2.1 Problem statement

This problem is based on a network that binds a large set of nodes $N = \{0, 1, \dots, n\}$, in which each node is in a different location. Node $N = 0$ is the starting point of the network, in this case, the warehouse or depot from which all products will originate. The rest of the nodes correspond to the stores that the retailer serves, also known as the customers of the network, and can be defined by $L = \{1, \dots, n\}$, a subset of N. The nodes are connected by edges $(i, j) \in E$, characterised by a time and a distance, measured in the corresponding units, and have an associated cost c_{ij}^k in monetary units (m. u.), linearly dependant on the distance between the nodes of the edge, and varying according to the type of vehicle *k* that performs route *r*.

These travels are performed by a fleet K of vehicles. Each vehicle $k \in K$ has an associated fixed cost fc_k and can only transport products of the category P_k, according to the temperature requisites of the product segment and the compatibility with the trailer used by the tractor.

In between these nodes travels a set of product segments P = {1, …, *p*} sent by the warehouse N = 0 and ordered by stores belonging to L. Let S = {(*l*, *p*): *l* ϵ L, *p* ϵ P} be the set of services *s* composed of a pair of a location *l* and a product segment *p*.

Since a product segment *p* encompasses various SKUs, it is important to know the demand for each SKU ordered by each store. Each service has a corresponding quantity (sum of all products belonging to the segment) demanded that can vary according to the day of the order. For that reason, d_s^w defines the quantity of service *s* ϵ S in scenario *w* ϵ Ω with a probability p_w of occurring. This happens because from the point of view of the logistics team it only matters the number of pallets of each segment that the store receives (and not the individual components of each pallet).

For this project, each scenario was considered to be a different day. Since the demand varies according to the scenario and the size of the fleet, it must be fixed throughout the different days, and a balance between having too many vehicles that won't be working for several days and having fewer vehicles that won't be able to fulfil demands in days of intense activity must be made. For that reason, the planner must choose, for each demand d_s^w, if the service must be delivered by a vehicle of the fleet or by a subcontracted vehicle that is more expensive than the fleet but can be used on an occasional basis.

If the reader intends to adapt this model to a case in which instead of having product segments there is a small number of SKUs to be transported and there is a possibility to have different delivery time windows per SKU then the set P can easily accommodate the individual SKUs and d_s^w will be the demand of each product, instead of the sum of all products belonging to a product segment,

To define a route *r* is important to know which subset of consecutive sequenced edges $E_r \subseteq E$ are traversed by the route and what services $S_r \subseteq S$ can be fulfilled in that journey considering all operational constraints, and what is the list $T_r \subseteq T$ of the arrival times to each location visited by the route. Each route has mandatory a specific vehicle k_r ϵ K associated with it and is also defined by three different timings: pt_r is the preparation starting time of the loading activities, st_r the route starting time and et_r the route ending time. The duration of the route dur_r and the route cost c_r are the final parameters needed to define the route.

At last, let G be the set of discrete time slots *g* ϵ G available within one scenario. Each route *r* and time slot *g* have a binary parameter e_{rgs} indicating if a service *s* is reached by route *r* during time slot *g*. Through this variable, the model will guarantee compliance with all operational constraints, which means that the arrival time of route *r* at the location *l* of service *s* is compatible with the endogenous window that begins in slot *g*. Furthermore, the existence of a variable h_{rg} indicating if a route *r* is active during a slot *g* will allow to count the number of vehicles active in all slots and limit this value to the dimension of the fleet. It is important to mention that if a service *s* only starts within one slot *g* per route *r*, the same route *r* will be active during various slots *g*.

To reduce the computational effort of this model, there was a group of characteristics of this problem that were used to filter the routes generated in the first phase namely:

- A large majority of stores prefer to receive the refrigerated categories, S^{fresh}, right before the opening of the store to ensure quality and fresh food is readily available for their clients, implying that the delivery window of these products has an additional constraint related to opening hours of the store.
- The preparation time pt_r is limited by the cut-off time t_p of each product segment *p* from which the pallets are ready to be loaded into the trucks.
- Large food retailers with stressed planning days might find the necessity to deliver products during closing hours of a store. For this, they must identify

which stores have the necessary characteristics ($a_s = 1$ if store *s* can receive autonomous deliveries) like their own stacker, a secure loading dock, available storage space and refrigerated chambers and so on. If so, a truck can be sent with the products during the night with a helper that can perform the unloading operations at an added cost. All routes with at least one autonomous delivery are reunited in subset Θ^{autom}.

Finally, even if not included in the current model, it is also important to mention the realities of the warehouse and drivers, which were included in a different way to, once again, reduce the computational effort.

To bring the warehouse operations into the model, it is necessary to incorporate the availability of loading docks during the day. Normally, the warehouses in food retail sectors use their docks to unload products from the suppliers and load them into trucks to send to stores during different periods of the day. They might even have docks that are closed during certain periods due to a lack of labour force. For that reason, it was important to ensure the existence of a parameter l_{rg} that manages if a route *r* is loading during time slot *g* or not, while $maxDocks_g$ indicates the total number of available docks at any given slot *g*.

In this sector, the planning department might also be responsible for managing the driver shifts. While a vehicle might be available for 24 hours a day, a driver might only drive a certain amount regulated by national labour laws. For that reason, and to ensure that at the moment of changing drivers, there is a small number of routes activated, it should be incorporated into the model a set H^{change} that contains the time slots *g* used for this purpose. As with the docks, $maxDrivers_g$ will ensure that the model doesn't use more drivers than the ones available at any slot *g*.

While the driving time is controlled by only using historical data, that complies with all existing regulations, the number of docks is only controlled in the validation phase.

The important sets used in the proposed methodology are summarized in Appendix A.

4.2.2 Variables and Parameters

This methodology can be defined as a two-stage stochastic optimization problem. For that, it demands the existence of a two-level decision, since while some variables are independent of the scenario, others will vary according to parameter *w*. The first level, or the assignment of a time window, uses $y_{sg} = \{0,1\}$ to assign the beginning of service *s* to time slot *g*, while the second level manages the operational vehicle routing. In this level, x_{rs}^w is a continuous variable that indicates the portion of service *s* that route *r* serves in scenario *w*. Second level binary variable $z_r^w = \{0,1\}$ represents if route *r* is active in scenario *w*. At last, the variable sub_s^w quantifies the portion of service *s* to be delivered by a subcontracted vehicle instead of by the retailers' fleet in scenario *w*.

Variable v_k^w represents the number of vehicles of type *k* used in scenario *w*. However it can be easily transfomed into a paremeter representing the number of vehicles of type *k* available.

So this proposed version of TWAVRPP is as follows in table 5.

Table 5 – Indexes, parameters and variables used in the optimization model

Indexes	
w	Scenario *w*
r	Route *r*
k	Vehicle of type *k*
s	Service *s*
g	Time slot *g*
l	Location *l*
p	Product segment *p*
Parameters	
p_w	Probability of occurrence of scenario *w*
c_r	Cost of performing route *r*
c_{sub}	Cost of subcontracting a vehicle
c_{change}	Cost of changing a window
fc_k	Fixed cost of a vehicle of type *k*
d_s^w	Demand for service *s* in scenario *w*
q_k	Capacity of the vehicle of type *k*
cap_{rk}	1 if the vehicle of type *k* performs route *r*, 0 otherwise
e_{rgs}	1 if route *r* can deliver service *s* at time slot *g*, 0 otherwise
g_{rgl}	1 if route *r* delivers to location *l* at time slot *g*, 0 otherwise
h_{rg}	1 if route *r* is active in time slot *g*, 0 otherwise
$serv_{sl}$	1 if service *s* belongs to location *l*, 0 otherwise
$delta_{sg}$	1 if service s starting at slot g differs from the original window by more than 30 minutes
$last_{rl}$	1 if location *l* is the last one visited by route *r*, 0 otherwise
Integer decision variables	
u_{rl}^w	1 if route *r* serves location *l* in scenario *w*, 0 otherwise
v_k^w	Number of vehicles of type *k* necessary in scenario *w*
y_{sg}	1 if time window of service *s* starts at slot *g*, 0 otherwise
z_r^w	1 if route *r* is performed in scenario *w*, 0 otherwise
Continuous Decision Variables	
x_{rs}^w	Portion of service *s* served by route *r* in scenario *w*
sub_s^w	Portion of service *s* subcontracted in scenario *w*

4.2.3 Objective Function

The original formulation of TWAVRPP intended to minimize the transportation costs of the operation, by minimizing the size of the necessary fleet and all travelling costs around all existing scenarios, represented by equation 4.3. The objective was to minimize the expected cost of performing all selected routes and the fleet required to do so.

$$minimize \sum_{w\in\Omega} p_w * \left(\sum_{r\in\Theta} c_r * z_r^w + \sum_{k\in K} fc_k * v_k^w\right) \tag{4.3}$$

In this project, there was an intention to incorporate into the objective function other relevant costs to the entire logistics operation, as is depicted by equation 4.4. The way to calculate the route cost c_r will differ significantly (as seen in equation 4.5) and a new cost of subcontracting part of the demand, c_{sub}, will be included. This might be a good solution to answer to abnormal peaks of demand without incrementing the size of the fleet. At last, a cost that doesn't depend on the scenario that accounts for the number of changes proposed by the model must also be added to the objective function. This cost will count the number of changes proposed that differ more than 30 minutes from the original window and multiply it by a cost c_{change}. In doing so, equation 4.4 intends to minimize, for all scenarios, the expected cost of delivering all products (through a vehicle of the fleet using a certain route or with a subcontracted vehicle) and the required fleet to do it, while also minimizing the number of changes done so that the model only changes windows that have a significant impact in the final value of the objective function.

$$minimize \left(\sum_{w\in\Omega} p_w * \left(\sum_{r\in\Theta} c_r * z_r^w + \sum_{s\in S} c_{sub} * sub_s^w + \sum_{k\in K} fc_k * v_k^w\right)\right) + \\ + \left(\sum_{s\in S}\sum_{g\in G} y_{sg} * delta_{sg}\right) * c_{change} \tag{4.4}$$

This formulation also allows studying a Pareto framework by changing the value of c_{change}. With a value of zero, it is possible to estimate the maximum operational savings if a retailer doesn't account for the difficulties of changing a window that is already defined. By incrementing the cost, the retailer can have a better understanding of how the costs depend on the number of changes made.

Furthermore, as previously mentioned, instead of only accounting for the travelling costs, there will be three other important costs to minimize. (4.5) defines the cost of a route as the sum of three different costs: c_{ij}^k, cost of traversing the edges of the network to execute the route; c_{inv}^r, investment cost to give the necessary conditions to a store; and c_{helper}^r, the cost of a helper if needed.

$$c_r = c_{inv}^r + c_{helper}^r + \sum_{k\in K}\sum_{(i,j)\in E_r} (c_{ij}^k * cap_{rk}) \ , \quad \forall r \in \Theta \tag{4.5}$$

With all three costs, the methodology intends to answer the biggest worries of all stakeholders. The cost of investment and the necessity of a new helper are costs incurred by moving a window from the day to the night period.

4.2.4 Constraints

The first set of constraints regulates the performance of all services that are required to deliver. Constraint (4.6) ensures that in all scenarios, all services *s* are satisfied, which means that in all scenarios *w* all product segments are delivered to all stores, either using the fleet or with a subcontracted service. (4.7) is a vehicle capacity requirement ensuring that there is no truck that needs to deliver more pallets than the ones that fit inside its trailer.

$$sub_s^w + \sum_{r \in \Theta} x_{rs}^w = 1 , \qquad \forall s \in S, w \in \Omega \tag{4.6}$$

$$\sum_{s \in S} d_s^w * x_{rs}^w \leq z_r^w * q_k * cap_{rk} , \qquad \forall r \in \Theta, w \in \Omega \tag{4.7}$$

Restriction (4.8) ensures that all services *s* have an associated endogenous time window *g* assigned to them, making sure that all services *s* are served, and with (4.9) the model only allows services to be satisfied by routes that arrive at the store inside said endogenous time window. Constraint (4.10) activates the routes that the model chooses to perform, and equation (4.11) will guarantee that a store is visited when a service is performed.

$$\sum_{g \in G} y_{sg} \geq 1 , \qquad \forall s \in S \tag{4.8}$$

$$x_{rs}^w \leq \sum_{g \in G} e_{rgs} * y_{sg} , \qquad \forall r \in \Theta, s \in S, w \in \Omega \tag{4.9}$$

$$\sum_{s \in S} x_{rs}^w \leq z_r^w , \qquad \forall r \in \Theta, w \in \Omega \tag{4.10}$$

$$u_{rl}^w \geq x_{rs}^w * serv_{sl} , \qquad \forall r \in \Theta, s \in S, w \in \Omega \tag{4.11}$$

To cover the limitations of the fleet, another constraint is needed. (4.12) counts the maximum number of vehicles of each type *k*, active at the same time in each scenario *w*. The last two equations intend to control two more operational realities of the deliveries. (4.13) only allows the model to activate one route per time slot *g* per store so that a store can only be visited by one vehicle per time slot. Equation (4.14) guarantees that all stores that need to see their respective tares collected are visited in the last place at least once by all active routes that pass through it.

$$\sum_{r \in \Theta} h_{rg} * z_r^w * cap_{rk} \leq v_k^w , \qquad \forall w \in \Omega, \mathrm{g} \in \mathrm{G}, \mathrm{k} \in \mathrm{K} \tag{4.12}$$

$$\sum_{r \in \Theta} u_{rl}^w * g_{rgl} \leq 1 , \qquad \forall\, l \in L, g \in G, w \in \Omega \tag{4.13}$$

$$\sum_{r \in \Theta} z_r^w * last_{rl} \geq 1 , \qquad \forall\, l, w \in \mathrm{Ltare}_w \tag{4.14}$$

(4.15) ensures that the sum of probabilities of all scenarios w is one and, at last, (4.16) limits the values of the decision variables according to the specifications previously mentioned.

$$\sum_{w\in\Omega} p_w = 1, \qquad \forall w \in \Omega \tag{4.15}$$

$$x_{rs}^{w},\ \ sub_s^w \in \mathbb{R}^+,\ u_{rl}^{w},\ \ y_{sg},\ \ z_r^w \in \{0{,}1\},\ v_k^w \in \mathbb{N}_0 \tag{4.16}$$

4.2.5 Additional constraints for future work

While the model used in the second phase ends in equation (4.16), there were also formulated another set of constraints to be used in future work that intend to model the filters applied in the first phase of the project and the realities related to the warehouse operations and drivers' workload.

Constraint (4.17) doesn't allow the delivery of services s via autonomous routes r to stores without the necessary conditions to receive the goods during closing hours.

$$x_{rs}^{w} \le cond_s \,, \forall r \in \Theta^{Auton} \,, s \in S, w \in \Omega \tag{4.17}$$

Equation (4.18) intends to guarantee that all selected routes respect the cut-off time of the products t_p of the services they carry, by limiting the preparation time pt_r of said routes according to the cut-off time at the warehouse.

$$e_{rgs} * z_r^w * l_{rg} * pt_r \ge t_p, \forall\, r \in \Theta, g \in W, w \in \Omega, s \in\ S, p \subset S(P) \tag{4.18}$$

(4.19) ensures that at any given moment g the number of trailers being loaded is never larger than the available number of docks and (4.20) restricts the necessity of drivers to the ones available for the operation at any moment g.

$$\sum_{r\in\Theta} l_{rg} * z_r^w \le maxDocks_g \,, \forall\, w \in \Omega, \mathrm{g} \in \mathrm{G} \tag{4.19}$$

$$\sum_{r\in\Theta} h_{rg} * z_r^w \le maxDrivers_g \,, \forall\, w \in \Omega, \mathrm{g} \in H^{change} \tag{4.20}$$

4.3 Phase 3: Validation

The final phase consisted of validating the results of the model with the three main stakeholders of the process.

As said in section 3.5, the project went through an iterative cycle between phases 2 and 3 to assure that the presented model strongly represented the real world.

The initial meetings were important to validate the model. Since the first iteration of the process wouldn't be perfect, it was important to ensure that all important assumptions and constraints were being included in the model, fix the ones that weren't correct, and add new ones that hadn't been initially mapped out. After that, it was important to understand how the model compared with the real planning that the planners had done and from such comparisons, new alterations had to be implemented.

To do so, it was important to validate the model for the current set of delivery time windows and check if the main KPIs originated from such runs were approximated to the ones of the planning team and check if the routes that were being selected by the model made sense for the planners.

Only then it was possible to conclude that the model was a faithful representation of the real world and only from that moment on the results obtained by the model were assumed to be correct and the degrees of freedom being given to the model could start to increase.

5 Results

The strategy developed to guarantee the model's validity was based on a comparative approach between a baseline scenario and the optimized scenario since the application of non-optimized time windows could generate millions of euros in losses to the retailer.

Firstly, the previously generated set of routes was filtered with each service's current delivery time windows, generating a subset of routes that are possible to be performed at the present, without making any changes. When applying the model to this subset, it was expected to obtain results very similar to the real world. The second step compares the baseline scenario with the real-world KPIs of the operation. If both were aligned, it was sufficient to prove the model was valid. Lastly, by removing the filter applied in the first phase, a new set of optimized time windows can be obtained, and the results easily compared with the previous baseline scenarios.

Such an approach is resumed in figure 9.

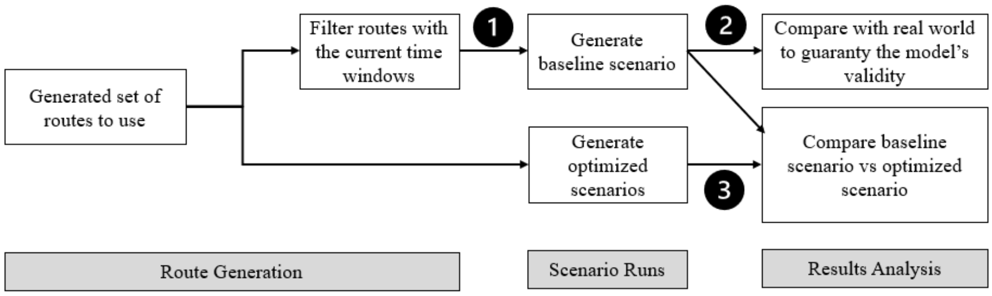

Figure 9 – Approach to validate the results

5.1 Instances Used

As was mentioned in Chapter 3, the retailer considers in its planning phase six different product segments – PBL, PBS, Fresh goods, Frozen goods, Meat and Fish – and owns more than 300 stores that are served by two warehouses (average of 150 stores per warehouse). For the present thesis, and since the project is still ongoing at the moment of the writing, only one of the warehouses was chosen to be analysed.

The partnership with one of the largest national food retailers gave us access to an entire year of data, so it was important to decide which data would be used to obtain the newly optimized set of routes.

The first thing important to analyse was the seasonality affecting the retailers' logistic operations throughout the year. For that, there were two important proxies to consider: the number of trucks used at the peak period of the day (namely the dawn) and the number of pallets delivered throughout the year, as shown in figure 10.

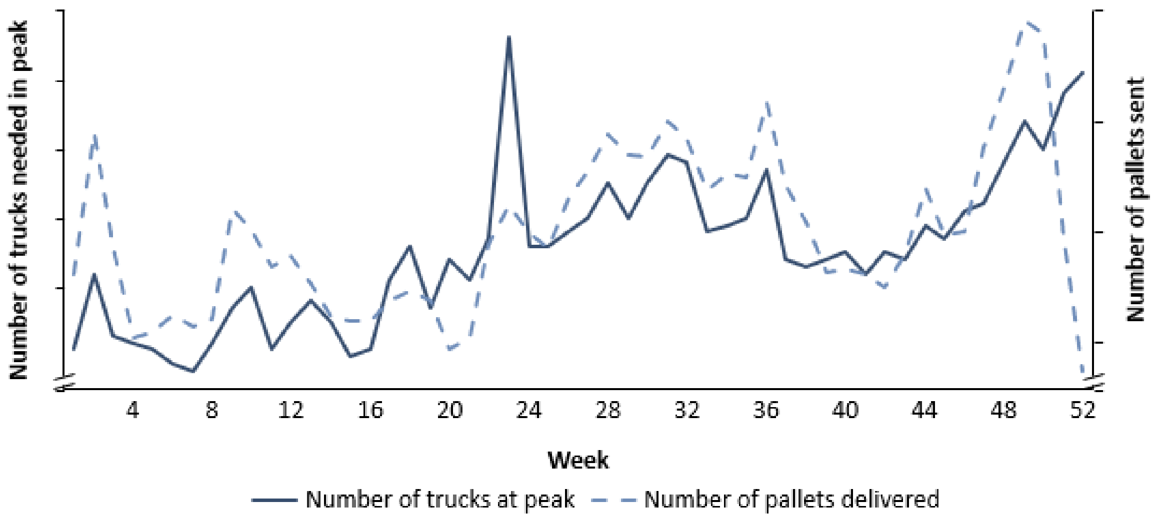

Figure 10 - Seasonality of the logistical operations from the point-of-view of 1) the fleet requirements at peak and 2) the number of pallets delivered

When looking at the number of trucks needed during the peak, figure 10 shows that the operation is more stressed in the second semester of the year with peak activity in the middle of May (week 22), a plateau during the summer period (between weeks 28 and 36) and a final increase in fleet requirements in the last weeks of the year near the Christmas holidays.

However, the second metric is important to understand if the increase in fleet requirements (which brings a large share of the operations costs) is related to a rise in demands or only to a combination of other factors that deemed the necessity of more vehicles, reducing the occupation of each one. With it, it is possible to observe that the peak during May is not due to demand, which cannot be said about the periods during the summer and the December holidays. Looking at the demand, it can also be observed that the peaks of early January and during March don’t influence the number of vehicles needed, which shows that during these periods, the fleet can be more easily optimized to deliver all pallets required with fewer trucks.

Moreover, the most important conclusion of this analysis is the importance of obtaining a robust set of optimized delivery time windows that can be the same throughout the entire year and can simultaneously answer two big concerns. Firstly, during peak periods of summer and Christmas holidays, the optimized set of windows must be complied with, without the necessity of increasing significantly the fleet required, and secondly, during lower periods, the fleet can be optimally used without having too many trucks stopped, incurring in fixed costs. From meetings with the retailer’s team, it came clear that it was important to avoid using weeks of extreme stress to the supply chain, like the summer in Algarve and the Christmas week, in which some windows and assumptions are not respected.

With all these things in mind, two weeks were chosen, according to the two metrics analysed: the first week within the percentile 50 and the second near the percentile 85. Figure 11 shows how the number of trucks needed varies according to the day of the week, and the hour of the day. Figure 12 corroborates this by showing that on those days when fewer trucks are needed, there is a smaller number of pallets being delivered. It also shows that Mondays and Sundays are days with almost half of the deliveries of the other days, as well as the differences between these two weeks. These fluctuations have an important impact on the delivery time windows since they must work for all days and all different demands.

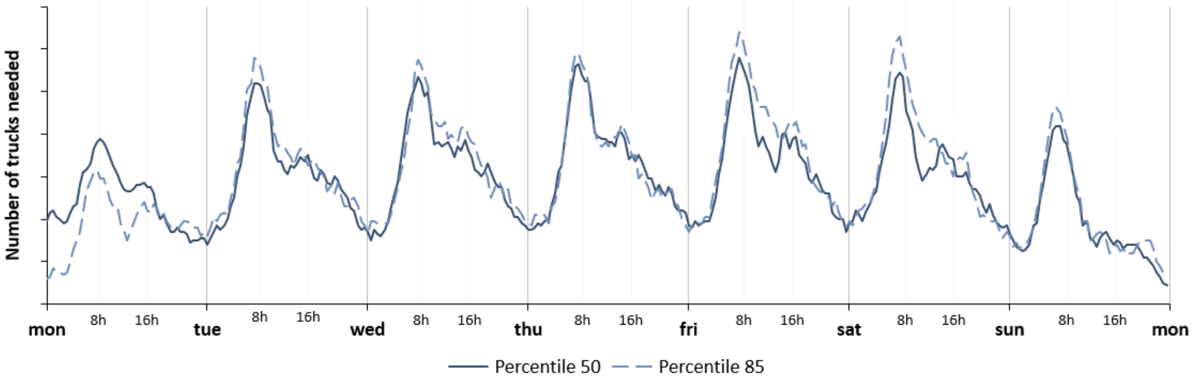

Figure 11 - Daily seasonality of the number of trucks needed

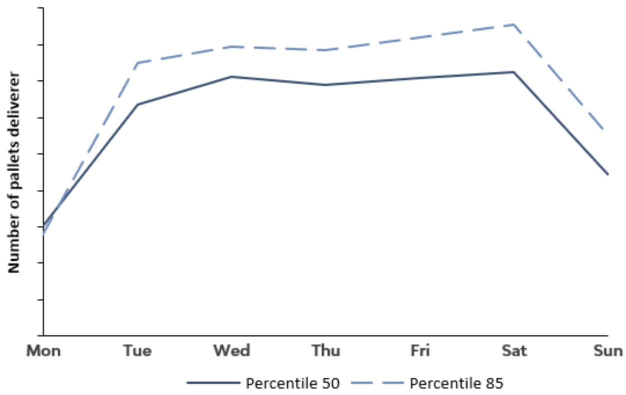

Figure 12 - Difference in the number of pallets being delivered on each day of both weeks

At last, figure 13 shows the variability of demands where the impact of both the ambient segments (PBS and PBL) and frozen goods is visible. On Mondays, there is a clear reduction in the deliveries of PBL and PBS segments, while the fish and meat deliveries are non-existent. This is counterbalanced by an increased number of deliveries of frozen goods that decrease significantly in the remaining days of the week. On Sundays, there are no deliveries of frozen goods and a clear reduction in the PBL sector. These dynamics allow for better management of the retailer's fleet and the reduction of the incurred costs.

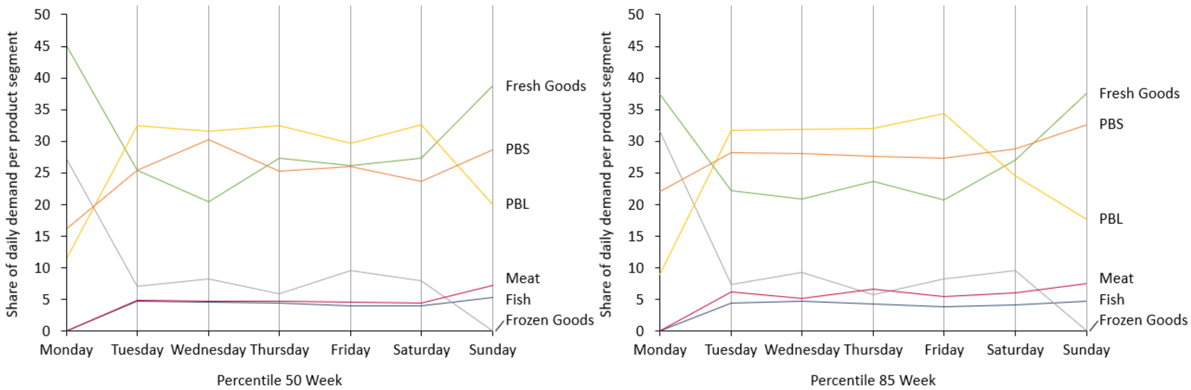

Figure 13 - Daily share of pallets sent, by product segment, in the chosen weeks

As previously mentioned, the two weeks chosen were used to create two scenarios. By filtering the set of routes and only allowing routes that comply with the current delivery time windows of each service, it is possible to generate the baseline scenarios. These scenarios will have two main objectives. Firstly, they allowed comparing the model with the real-world KPIs and the model's validation. Secondly, they will be used to compare with the results of the optimized scenarios. By analysing table 6 it is easy to understand that the difference in demands mentioned previously is responsible for a significant increase in the operational costs of the most stressed days (Friday and Saturday) and in the fleet requirements: both in terms of the total number of vehicles required at peak and in the number of vehicles being loaded at the same time.

Moreover, the increase in demand implies a significant increase in subcontracting vehicles outside the current fleet, which will also helps in rising the operational cost of the day. At last, the increase in the travelled distance is also explained by the rising in demand that can generate a necessity to send more trucks to stores far from the warehouse, since the demand might not fit into the number of trucks that fitted previously.

Table 6 - KPIs variation between baseline instances for both weeks

KPI	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday
Total Cost (€)	-0.48%	8.84%	14.58%	29.00%	40.85%	40.10%	32.46%
Travelled distance (km)	2.88%	12.58%	4.79%	4.09%	2.84%	14.45%	3.62%
Number of autonomous windows	0.00%	0.00%	0.00%	25.00%	16.67%	0.00%	-16.67%
Pallets delivered with fleet	-0.58%	16.60%	13.75%	14.26%	14.56%	16.78%	16.02%
Pallets delivered via subcontracted service	0.00%	-3.95%	92.54%	51.85%	100.00%	101.35%	88.46%
Vehicles required at peak	-5.00%	5.95%	1.22%	12.05%	5.75%	16.28%	6.49%
Maximum number of vehicles in preparation	0.00%	14.29%	-2.08%	18.60%	8.33%	15.22%	6.52%

After choosing the weeks to study, filtering the set of routes with the current delivery time windows, and running the model to obtain the KPIs previously mentioned, it was important to compare the results obtained with the ones observed by the retailer on those specific days. One of the most valuable analyses was to compare the routes selected by the model with those executed by the retailer's team and understand if the number of vehicles needed by the model was similar to the ones used. Figure 14 shows that comparison where we can see a close resemblance between both lines. The differences observed were deeply analysed and considered insufficient to invalidate the model. The main cause of those differences was the selection of some different routes to fulfil the same services.

All fourteen scenarios were obtained by running the optimization model for one hour, which was not enough to get a gap of zero (shown in table 7). This problem was more significant for the percentile 85 week, especially for the days when the supply chain is more stressed due to an increase in demand. It is important to note that in MIP problems the gap measures the difference between the value of the objective function of the incumbent solution and a possible best bound and, for that reason, even if a gap different than 0 can be observed in an optimal solution, a gap close to 0 gives more confidence in that the optimal solution was found.

Table 7 – GAP of the baseline instances

GAP	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday
Percentile 50	0.88%	7.12%	8.24%	5.98%	11.94%	14.43%	1.86%
Percentile 85	1.00%	11.47%	9.10%	19.85%	20.59%	19.44%	1.77%

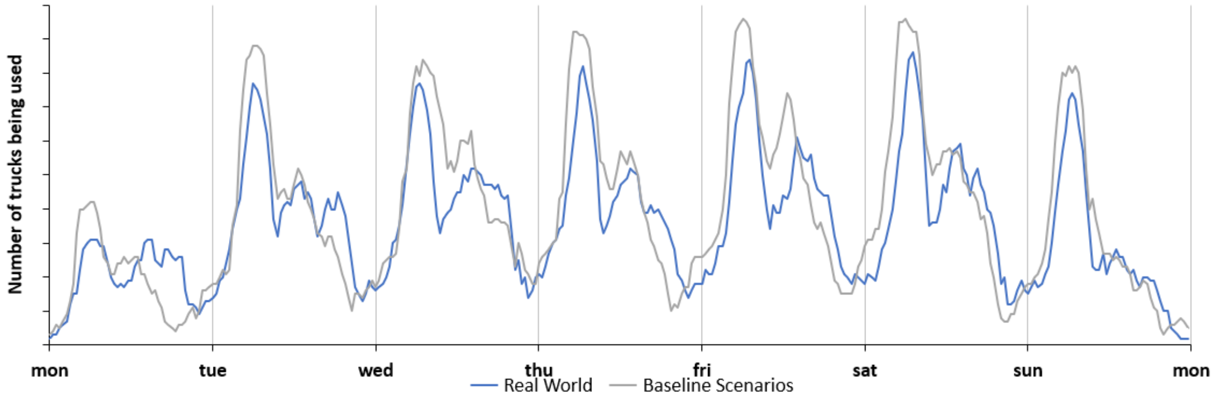

Figure 14 - Comparison between the number of vehicles used by the model and the ones used by the retailer during the percentile 50 week

5.2 Optimized Results

After having the model validated and the main KPIs of the baseline scenarios aligned with the retailer's values, it was possible to start using the model to optimize a new set of delivery windows. The comparison between the main KPIs of both weeks and the respective baselines will be analysed in this section but can be found in more detail in Appendix B.

Before the analysis, it is important to mention that since the project is still ongoing at the time of the writing of this thesis, the results presented are still preliminary.

The first important thing to review is the number of proposed changes by the optimization model to the current delivery windows. A first look at those changes allowed us to find some changes of less than 30 minutes that were only small corrections to windows that did not start in one of the slots *g,* and so shouldn't be included. For example, a window starting at 07:15 was changed to begin at 07:00. After excluding alterations of less than 30 minutes, the obtained results were condensed in Table 8. It is important to note that these are the only changes that will be accounted for in the objective function.

Table 8 - Percentage of alterations proposed to the delivery time windows

Type of Change	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday
Percentile 50 week							
Unaltered	74%	75%	77%	75%	73%	78%	74%
Anticipation	17%	16%	16%	17%	17%	14%	18%
Delayed	9%	9%	8%	8%	9%	8%	8%
Percentile 85 week							
Unaltered	77%	69%	86%	69%	77%	71%	74%
Anticipation	14%	24%	2%	23%	13%	17%	19%
Delayed	9%	7%	12%	8%	10%	12%	7%

As mentioned in section 3.3, a store manager would only accept having a single window per segment. For the final recommendation, the windows proposed for each day should be cross-checked and later re-tested to see the impact of the final windows decided upon. Table 8 shows that to reduce operational costs, the model prefers to anticipate deliveries than delay them. Furthermore, it is easy to see that, on average, there were more changes in the second week than in the first, which is also another indicator that a multiday instance should be used, to optimize the delivery windows for several days at the same time.

Such changes to the windows can be further observed in figure 15 in which the number of pallets delivered on the Thursday instances (percentile 50 week) with both sets of windows (the original and the optimized) is represented. Lines A and B help to analyse the delivery peaks for both instances, while line C tries to show the behaviour of the deliveries throughout the rest of the day.

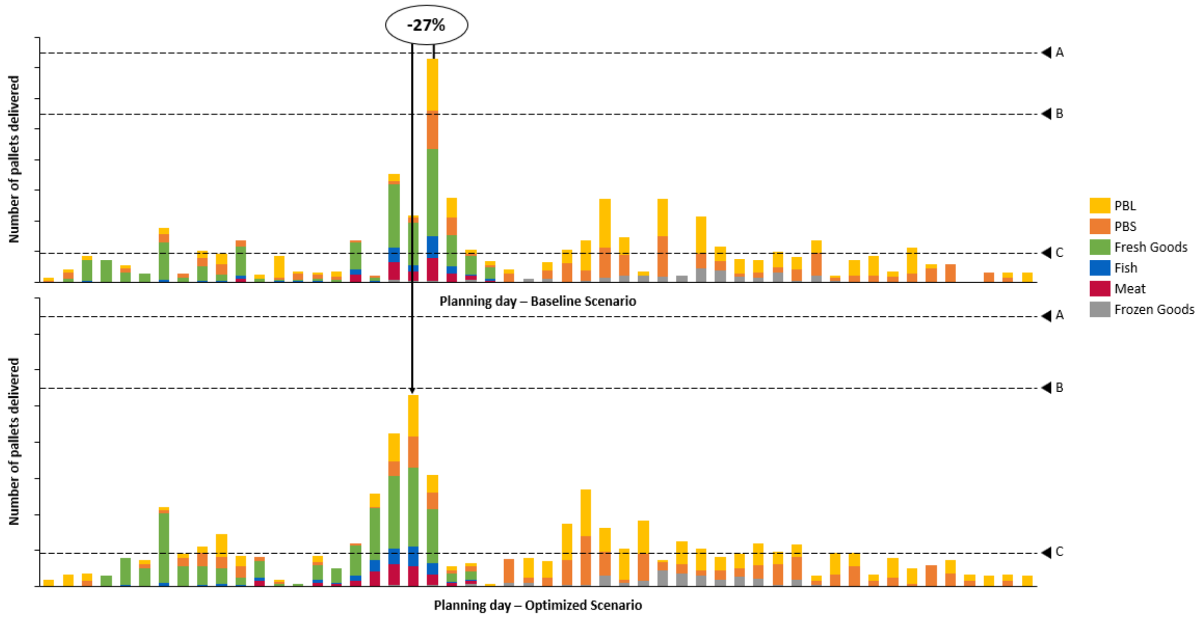

Figure 15 - Comparison of the number of pallets delivered throughout a day of the percentile 50 week with the original set of windows and the optimized ones

There is a clear reduction in the delivery peak (from A to B) of 27%, which will reduce the congestion inside the warehouse, the number of trucks needed to deliver all those pallets and an increment in deliveries throughout the day (more columns of the graphic closer to line C). It is also possible to verify that all fresh products are delivered at the beginning of the day to assure that they are already in the store when they open. In contrast, the rest of the product segments are distributed mainly during the rest of the day.

Since the main objective of this process was to optimize the delivery windows to reduce the logistical costs of the operator, is it important to look at the behaviour of the objective function for both sets of instances. When looking at the operational costs savings of each instance, present in table 9, it can be seen daily savings of more than 50% of the baseline for both weeks under analysis.

Table 9 - Estimated savings for the optimized set of time windows

	Monday	**Tuesday**	**Wednesday**	**Thursday**	**Friday**	**Saturday**	**Sunday**
Percentile 50 week	53.78%	62.21%	57.98%	57.04%	59.95%	58.88%	59.98%
Percentile 85 week	53.59%	57.47%	59.70%	58.62%	69.39%	65.34%	65.91%

Due to the higher costs of subcontracting services, it is easy to understand that in the baseline instances this is the item of the objective function with a higher share of the costs. Meanwhile, with the optimization of the windows, the model will try to reduce this portion of the operational costs. Figure 16 shows how the share of subcontracted costs decreases significantly from the baseline instances to the optimized scenarios, while not incurring in increments on the other rubrics of the objective function.

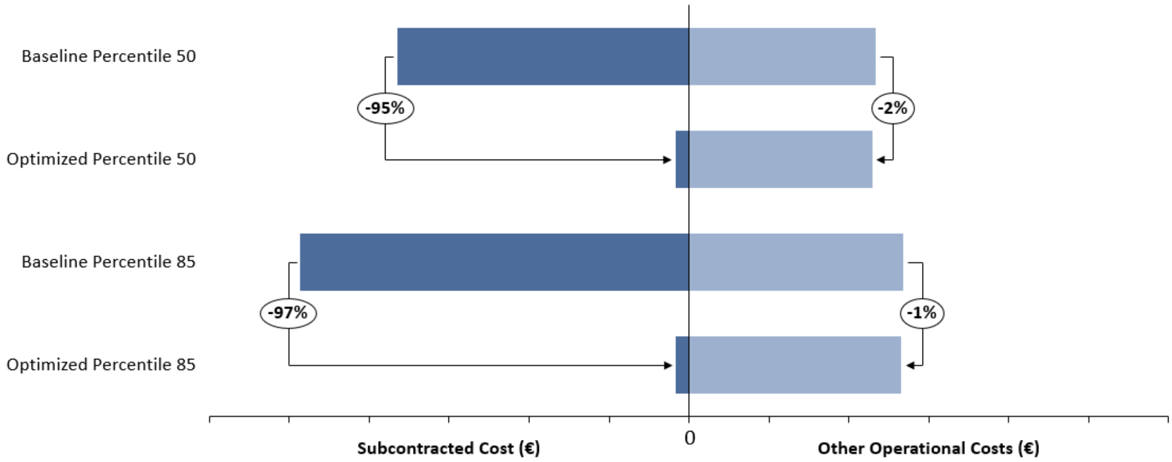

Figure 16 - Comparison between the variation of subcontracted costs and other rubrics for the baseline and optimized instances of both weeks under analysis

A deeper analysis of the results will further corroborate the values presented. By optimizing the windows, two things were expected: 1) reduction of the required fleet and 2) reduction of the subcontracted travels or, if possible, not even require them. Table 10 summarises the daily reduction in the number of trucks required for each scenario. As expected, days with less stressed planning, like Mondays and Sundays, allow for bigger reductions on the fleet, while in the remaining week, such effects are less observed.

Table 10 - Estimated reduction of vehicles needed during peak

	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday
Percentile 50 week	20.00%	9.52%	9.76%	1.20%	9.20%	8.14%	10.39%
Percentile 85 week	15.79%	12.36%	4.82%	7.53%	11.96%	5.00%	6.10%

The simplest analysis of table 10 shows that on all days the number of required vehicles is reduced. However, reality taught us another way to look at these values. Since the retailer doesn't rent vehicles daily, it is more interesting to look at the maximum value used during the entire week. Since the retailer is paying for all vehicles in the fleet, it is more interesting to use them all than to have them parked and not get some rentability out of them (maximizing both the occupation and the utilization rates of the fleet). If we analyse figure 17, we can see that even from this perspective, we can guarantee a reduction of 6% on the maximum number of vehicles needed throughout the entire week.

Finally, when analysing the effect of the new windows on the fleet requirements, the impact on the preparation phase is not very significant, with both the baseline scenario and the optimized scenario having a similar behaviour related to the number of trucks being loaded in the warehouse docks. This means that the optimized model doesn't increase the traffic inside the warehouse, ensuring a small disruption in the current operations. Such behaviour is best described in figure 18.

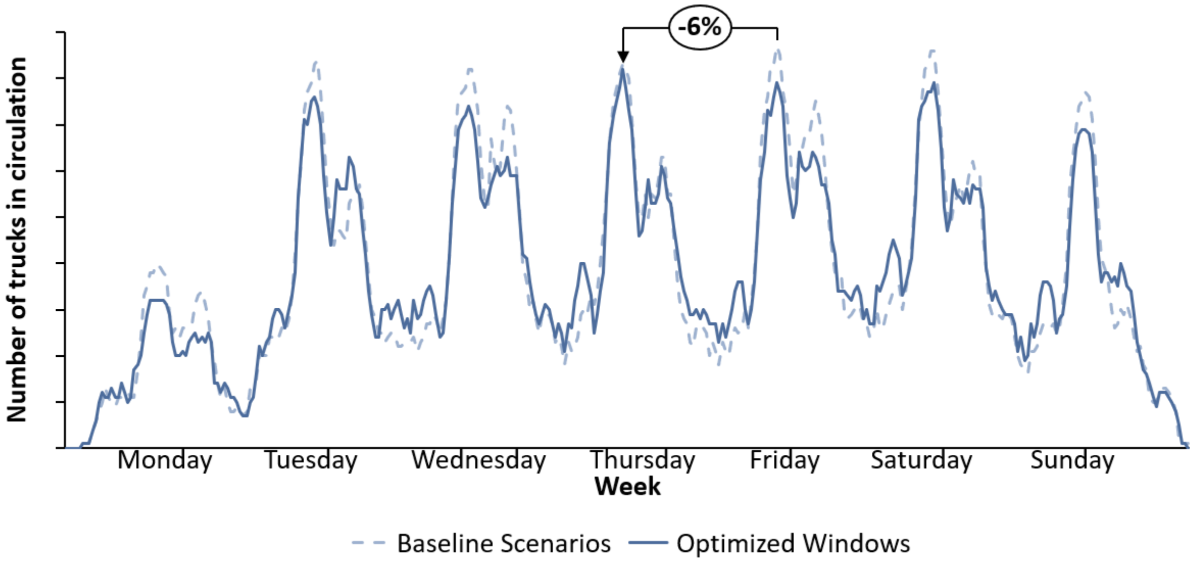

Figure 17 - Number of vehicles in circulation throughout an entire week (percentile 50) for both scenarios

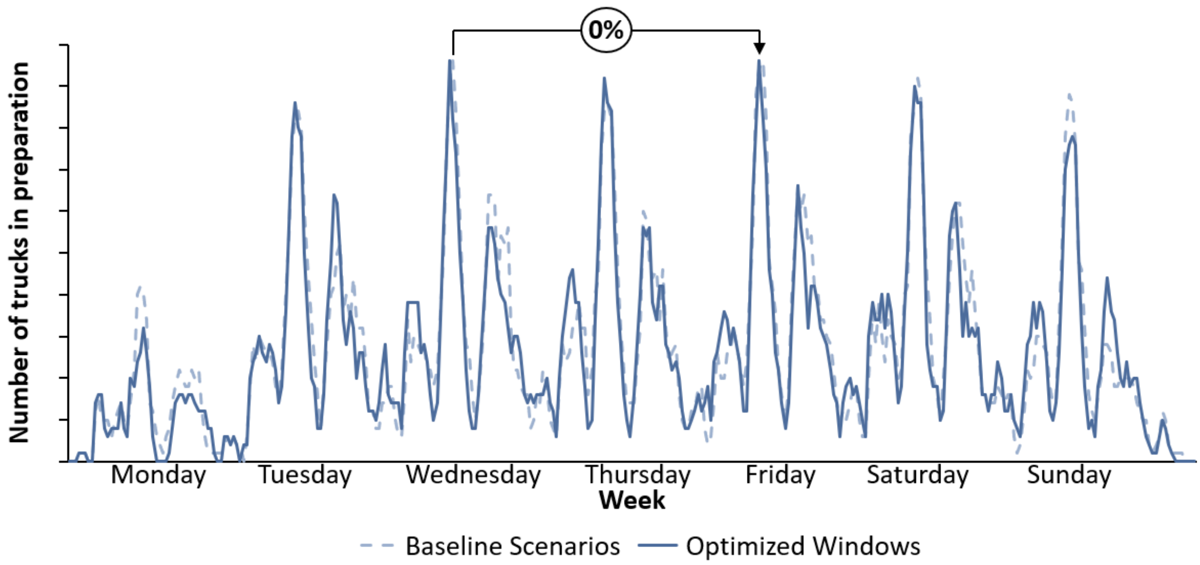

Figure 18 - Number of vehicles being prepared at the warehouse throughout an entire week (percentile 50) for both scenarios

Furthermore, it might be important to analyse the impact of the changes in specific parts of the warehouse operation. As was mentioned in section 3.1, the retailer uses different sections of the warehouse complex to receive, handle and dispatch each product segment. For that reason, even if the analysis of figure 18 shows an overall alignment between the number of vehicles being loaded in the optimized scenarios compared to the baselines, there might be a hidden effect on the warehouse operations of specific segments. As mentioned in section 4.2.1 the number of docks available was not included in the optimization model but was analysed in the validation phase. The warehouse team defined that only the three segments shown in figure 19 needed to be analysed since they corresponded to the more stressed operations, due to the higher volumes expedited daily.

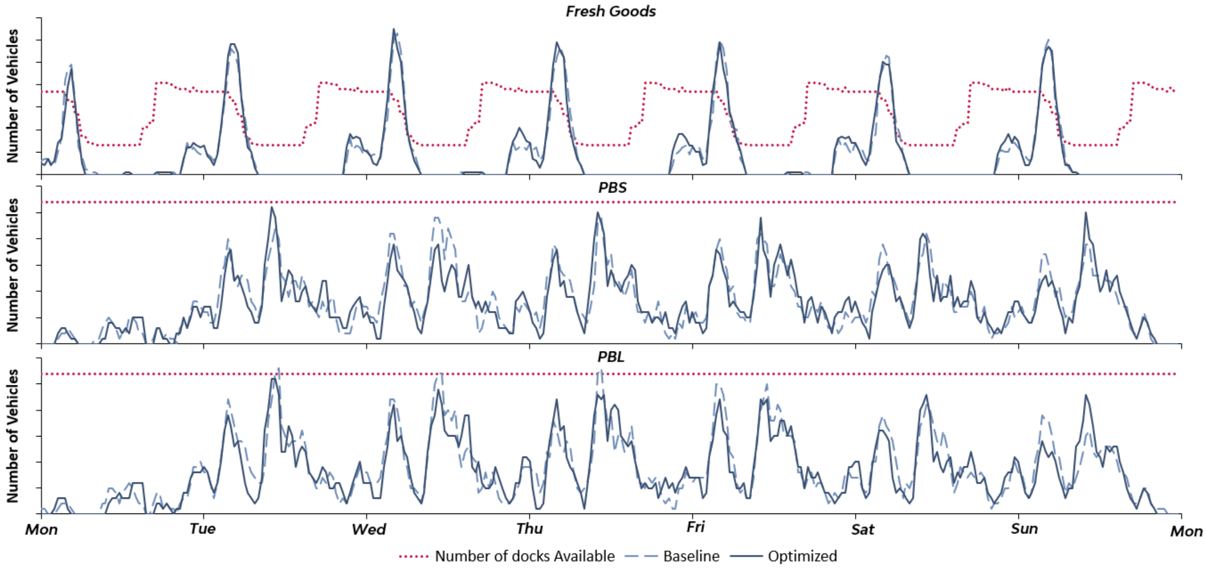

Figure 19 - Comparison between the number of docks available and the number of vehicles being prepared

While the windows' optimization decreased the loading peaks in both PBS and PBL, the number of trucks far exceeded the number of available docks to load the pallets of Fresh goods. The comparison with the baseline scenarios showed an alignment between both and further reunions with the team. This allowed to conclude that the anticipation of loads that are currently done by the warehouse team is sufficient to answer to such peaks. And so, the impact on the warehouse operations is corroborated to be insignificant, if not positive.

Another source of the cost reduction is related to the subcontracted vehicles being used. Since the optimized set of windows eliminates this type of vehicle (as seen in table 11), the cost associated with them will also be eliminated.

Table 11 - Share of pallets delivered via subcontracted vehicles on both weeks

	Monday	**Tuesday**	**Wednesday**	**Thursday**	**Friday**	**Saturday**	**Sunday**
Percentile 50 week							
Baseline	0.64%	1.48%	1.16%	0.99%	1.44%	1.21%	0.65%
Optimized	0.00%	0.00%	0.02%	0.00%	0.00%	0.02%	0.00%
Percentile 80 week							
Baseline	0.64%	1.22%	1.95%	1.31%	2.48%	2.06%	1.05%
Optimized	0.00%	0.00%	0.06%	0.00%	0.00%	0.03%	0.00%

6 Conclusions and future work

The main objective of this dissertation was to study the best way to change the delivery time windows of multiple product segments to the stores of large food retailers, minimizing not only the transportation costs but all the logistical costs of the operation. Since the retailers already have a set of time windows assigned to each store, this problem must also tackle the entropy that changing each window will cause.

An extensive literature review of the VRP was done to understand what was the best model to use to solve the problem at hand. Several versions of PVRP, ConVRP and VRPTW were studied, as well as the versions of TWAVRP known to the author. From all of them, it was easy to conclude that the best version was the TWAVRPP since it is the only known model to tackle product-dependent variables to assign different delivery time windows.

The proposed methodology is based on an approach that intends to find near-optimal solutions based on a previously filtered subset of feasible routes.

To validate the proposed model, it was used in a real-world case study of one of the largest national food retailers with more than 300 stores that received 6 different product segments daily at different times of the day.

The method presented in this thesis starts with a route generation phase that aims to generate a large set of routes that will be the input of the optimization model. Since there was access to a year of performed routes by the retailer, this method was chosen to the detriment of others proposed by the literature like the LNS or nearest neighbour heuristic. The historical data only needed to be unfolded throughout the day to allow that any service could be fulfilled at any time. This phase was also important to help reduce the computational effort of the optimization model. To do so, a group of filters was used to ensure that all routes that were inserted into the model were feasible and complied with all existing operational constraints. This procedure decreased the number of constraints that the model was subject to. This thesis also presents the database structure used to store and manage all the used data.

Next, a two-stage stochastic optimization model was developed with two main objectives. The model assigns time windows to each service, independently of the scenario being used, and allocates each service in each scenario to a route so that a routing plan is made. For future developments, some constraints were presented that weren’t included in this phase but can be easily added to the model.

The third phase of the project, and maybe the most important one, was validating the model to ensure that it was a strong representation of the reality lived by the retailer. This phase also allowed us to identify problems with the two precedent phases and correct them.

To validate the model an analysis of the yearly deliveries and of the used vehicles allowed us to choose two weeks that represent a satisfying picture of the weekly demand variability. To obtain the baseline scenarios, the set of routes generated in phase 1 of the methodology was filtered using the current delivery windows of the retailer, so that a service is never visited outside its window.

When the model was applied to this set of filtered routes, it suggested the current windows as the optimized ones since there were no degrees of freedom to change them. Then the KPIs of each instance were calculated and compared with the KPIs measured by the retailer's team on those specific days. After improving the model, the results were close enough, and the model was validated. From this point on, the current delivery window filtered applied to the routes was lifted, and it was possible to start optimizing a new set of windows.

Some comparisons between the baseline scenarios and the optimized ones are presented. The preliminary results allowed us to obtain an operational cost reduction of more than 50% by changing around 25% of the current windows.

These changes decreased the delivery peak by 30%, leading to a reduction of at least 6% of the vehicles required during the entire week, with days where a large reduction of vehicles can be seen if new contracts can be redacted for the vehicle rental schemes. The impact of these changes on the warehouse operations was also studied, concluding that the number of trucks being loaded at any time remained very similar.

The operational costs are also reduced by eliminating the subcontracted vehicles used by the baseline instances since these vehicles have a larger cost than the normal fleet.

Since the project is still ongoing during the writing of this thesis, the results presented are for individual daily instances, and not for a larger set of scenarios at the same time, and some still present a value of gap that is not acceptable for the final recommendation.

For that reason, there is a clear path still to follow. The model should be run with a larger run time so that the model tries to reduce the gaps presented. Secondly, instances with more than an individual day should be run so that the proposed set of optimized windows doesn't vary according to the day, and the results allow for a better understanding of the impact of the changes on the retailers' operational costs.

Furthermore, it is possible to understand that creating instances with more than one day will significantly increase the computational effort of the model and, for that reason, a Fix-and-Optimize matheuristic should be developed to help in this regard.

Another advantage of the proposed methodology is the existence of a parameter in the objective function that can be changed to study the Pareto framework of how the operational costs vary according to the number of changes and what is the impact of implementing these changes in the operation.

Regarding how to develop further the work presented in this thesis, some clear alternatives exist.

About the first phase of route generation, some of the heuristics used in the literature can be included to generate more diverse routes, such as nearest neighbour or ALNS. There might also be interesting to perform a robust analysis of the results obtained by routes generated by these methods when compared with those obtained with the historical routes as proposed in this thesis and understand which would be the best method to follow.

Another possible path relates to the warehouse stakeholder of this operation. The integration of a simulation model into the methodology could provide a deep study into the impact of the new proposed time windows on warehouse operations and allow for a more robust solution. This avenue would also help quantify the workforce needed for such operations.

A second optimization model or modifying the existing one could be an added value to this methodology, to allocate each performed route selected to a specific 12-hour shift of each vehicle, to create a clearer look into the daily plan associated with the proposed set of time windows. This would allow for a better dimensioning of the fleet and a more accurate analysis of the cost incurred with each solution.

It is possible to conclude that the 3 phased hybrid methodology is a good strategy to solve the TWAVRPP. The mathematical formulation allows for a close representation of the reality that the retailers face every day. The validation phase is essential to guarantee the quality of the results and conclusions presented.

By only using historical routes, the model guarantees that the stakeholder can more easily accept the solution and that a large set of constraints is already being complied with. By adding more costs to the cost of performing a route, the model comes closer to evaluating the impact of performing a route at a different time. At the same time, the changes proposed to the objective function are important to reduce changes that do not bring benefits to the operational cost but brings complications when negotiating such changes.

The preliminary results seem very promising and the identified next steps to improve the model will bring even more confidence in the validity of the model and its usefulness for future implementation in other case studies, in a world that is fast-moving with more demanding customers and stressed supply chains.

References

Portugal Grocery Retailer Preference Index: Second Edition. Dunnhumby, (2020). Last accessed in June 2022: https://www.dunnhumby.com/wp-content/uploads/2020/11/dunnhumby-Portugal-Grocery-RPI-Second-Edition.pdf

Baldacci, R., Bartolini, E., Mingozzi, A., and Valletta, A. (2011). An Exact Algorithm for the Period Routing Problem. *Operations Research* 59 (2): 228–41. https://doi.org/10.1287/opre.1100.0875.

Baldacci, R., Mingozzi, A., and Roberti, R. (2012). Recent exact algorithms for solving the vehicle routing problem under capacity and time window constraints. *European Journal of Operational Research* 218 (4): 1–6. https://doi.org/10.1016/j.ejor.2011.07.037.

Christofides, N., and Beasley, J. E. (1984). The period routing problem. *Networks* 14 (6): 237–56. https://doi.org/10.1002/net.3230140205.

Cordeau, J.-F., Laporte, G., and Mercier, A. (2001). A Unified Tabu Search Heuristic for Vehicle Routing Problems with Time Windows. *The Journal of the Operational Research Society* 52: 928–36. http://www.jstor.org/stable/822953.

Dalmeijer, K., and Spliet, R. (2018). A branch-and-cut algorithm for the Time Window Assignment Vehicle Routing Problem. *Computers and Operations Research* 89. https://doi.org/10.1016/j.cor.2017.08.015.

Dantzig, G. B., and Ramser, J. H. (1959). The Truck Dispatching Problem. *Management Science* 6: 80–91. http://www.jstor.org/stable/2627477.

Errico, F., Desaulniers, G., Gendreau, M., Rei, W., and Rousseau, L.-M- (2016). A priori optimization with recourse for the vehicle routing problem with hard time windows and stochastic service times. *European Journal of Operational Research* 249 (2): 55–66. https://doi.org/10.1016/j.ejor.2015.07.027.

Gaudioso, M., and Paletta, G. (1992). A Heuristic for the Periodic Vehicle Routing Problem. *Transportation Science* 26 (1992): 86–92. http://www.jstor.org/stable/25768524.

Groër, C, Golden, B., and Wasil, E. (2009). The Consistent Vehicle Routing Problem. *Manufacturing & Service Operations Management* (10). https://doi.org/10.1287/msom.1080.0243.

Gutierrez, A., Dieulle, L., Labadie, N., and Velasco, N. (2016). A multi population memetic algorithm for the vehicle routing problem with time windows and stochastic travel and service times. *IFAC-PapersOnLine* 49: 1204–9. https://doi.org/10.1016/j.ifacol.2016.07.673.

Helber, S., and Florian Sahling, F. (2010). A fix-and-optimize approach for the multi-level capacitated lot sizing problem. *International Journal of Production Economics* 123 (2): 247–56. https://doi.org/10.1016/j.ijpe.2009.08.022.

Ichoua, S., Gendreau, M., and Potvin, J.-Y. (2003). Vehicle dispatching with time-dependent travel times. *European Journal of Operational Research* 144 (1): 379–96. https://doi.org/10.1016/S0377-2217(02)00147-9.

Ioachim, I, Gélinas, S., Soumis, F., and Desrosiers, J. (1998). A dynamic programming algorithm for the shortest path problem with time windows and linear node costs. *Networks* 31 (5): 193–204. https://doi.org/10.1002/(SICI)1097-0037(199805)31:3<193::AID-NET6>3.0.CO;2-A.

Jalilvand, M, Bashiri, M., and Nikzad. E. (2021). An effective Progressive Hedging algorithm for the two-layers time window assignment vehicle routing problem in a stochastic environment. *Expert Systems with Applications* 165 (3): 113877. https://doi.org/10.1016/j.eswa.2020.113877.

Kallehauge, B. (2008). Formulations and exact algorithms for the vehicle routing problem with time windows. *Computers & Operations Research* 35 (7): 2307–30. https://doi.org/10.1016/j.cor.2006.11.006.

Kovacs, A. A., Golden, B. L., Hartl, R. F., and Parragh, S. N. (2015). 'The Generalized Consistent Vehicle Routing Problem'. *Transportation Science* 49: 796–816. http://www.jstor.org/stable/43666921.

Li, F., Golden, B., and Wasil, E. (2005). Very large-scale vehicle routing: new test problems, algorithms, and results. *Computers & Operations Research* 32 (5): 1165–79. https://doi.org/10.1016/j.cor.2003.10.002.

Li, M., Lu, J., and Ma, W. (2019). A hybrid multi-objective genetic algorithm for bi-objective time window assignment vehicle routing problem.

Martins, S., Ostermeier, M. Amorim, P., Hübner, A., and Almada-Lobo, B. (2019). Product-oriented time window assignment for a multi-compartment vehicle routing problem. *European Journal of Operational Research* 276. https://doi.org/10.1016/j.ejor.2019.01.053.

Neves-Moreira, F., da Silva, D. P., Guimarães, L., Amorim, P., and Almada-Lobo, B. (2018). The time window assignment vehicle routing problem with product dependent deliveries. *Transportation Research Part E: Logistics and Transportation Review*. https://doi.org/10.1016/j.tre.2018.03.004.

Pisinger, D., and Ropke, S. (2007). A general heuristic for vehicle routing problems. *Computers & Operations Research* 34 (8): 2403–35. https://doi.org/10.1016/j.cor.2005.09.012.

Ridder, R. J. (2014) ConRTR, a Consistent Vehicle Routing Problem algorithm.

Solomon, M. (1987). Algorithms for the Vehicle Routing and Scheduling Problems with Time Window Constraints. *Operations Research* 35: 254–65. http://www.jstor.org/stable/170697.

Spliet, R., and Gabor, A. F. (2015). The time window assignment vehicle routing problem. *Transportation Science* 49 (11): 721–31. https://doi.org/10.1287/trsc.2013.0510.

Spliet, R., Dabia, S., and van Woensel, T. (2018). The time window assignment vehicle routing problem with time-dependent travel times. *Transportation Science* 52 (2018): 261–76. https://doi.org/10.1287/trsc.2016.0705.

Spliet, R., and Desaulniers, G. (2015). The discrete time window assignment vehicle routing problem. *European Journal of Operational Research* 244. https://doi.org/10.1016/j.ejor.2015.01.020.

Stavropoulou, F., Repoussis, P. P., and Tarantilis, C.D. (2019). The Vehicle Routing Problem with Profits and consistency constraints. *European Journal of Operational Research* 274 (4): 340–56. https://doi.org/10.1016/j.ejor.2018.09.046.

Subramanyam, A., Wang, A., and Gounaris, C. (2018). A scenario decomposition algorithm for strategic time window assignment vehicle routing problems. *Transportation Research Part B: Methodological* 117 (11): 296–317. https://doi.org/10.1016/j.trb.2018.09.008.

Subramanyam, A., and Gounaris, C. (2016). A branch-and-cut framework for the consistent traveling salesman problem. *European Journal of Operational Research* 248 (1): 384–95. https://doi.org/10.1016/j.ejor.2015.07.030.

Subramanyam, A., and Gounaris, C. (2017). A Decomposition Algorithm for the Consistent Traveling Salesman Problem with Vehicle Idling. *Transportation Science* 52 (4). https://doi.org/10.1287/trsc.2017.0741.

Taillard, E., Badeau, P., Gendreau, M., Guertin, F., and Potvin, J.-Y. (1997). A Tabu Search Heuristic for the Vehicle Routing Problem with Soft Time Windows. *Transportation Science* 31 (5): 170–86. https://doi.org/10.1287/trsc.31.2.170.

Tarantilis, C. D., Stavropoulou, F., and Repoussis, P. P. (2012). A template-based Tabu Search algorithm for the Consistent Vehicle Routing Problem. *Expert Systems with Applications* 39 (3): 4233–39. https://doi.org/10.1016/j.eswa.2011.09.111.

Taş, D., Dellaert, N., van Woensel, T., and de Kok, T. (2013). Vehicle routing problem with stochastic travel times including soft time windows and service costs. *Computers & Operations Research* (1). https://doi.org/10.1016/j.cor.2012.06.008.

Taş, D., Dellaert, N., van Woensel, T., and de Kok, T. (2014a). The time-dependent vehicle routing problem with soft time windows and stochastic travel times. *Transportation Research Part C: Emerging Technologies* (11). https://doi.org/10.1016/j.trc.2014.08.007.

Taş, D., Jabali, O., and Van Woensel, T. (2014b). A Vehicle Routing Problem with Flexible Time Windows. *Computers & Operations Research* (12). https://doi.org/10.1016/j.cor.2014.07.005.

Vareias, A. D., Repoussis, P. P., and Tarantilis, C. D. (2019). Assessing customer service reliability in route planning with self-imposed time windows and stochastic travel times. *Transportation Science* 53. https://doi.org/10.1287/trsc.2017.0748.

Wang, Y., Zhang, S., Guan, X., Peng, S., Wang, H., Liu, Y., and Xu, M. (2020). Collaborative multi-depot logistics network design with time window assignment. *Expert Systems with Applications* 140 (2). https://doi.org/10.1016/j.eswa.2019.112910.

Xu, Z., and Cai, Y. (2018). Variable neighborhood search for consistent vehicle routing problem. *Expert Systems with Applications* 113 (12): 66–76. https://doi.org/10.1016/j.eswa.2018.07.007.

APPENDIX A: Important sets used in the problem formulation

P	Set of product segments p
L	Set of locations l
$Ltare_w$	Subset of location l who require collecting tare in scenario w
S	Set of services s, pair of a product h and a place l: $S = \{(l, h): l \in L, h \in P\}$
K	Set of vehicles k
Ω	Set of scenarios w
G	Set of discrete timeslots g available at each scenario w
Θ	Set of routes r

APPENDIX B: KPIs comparison between baseline and optimized scenarios

Table B.1 - KPIs comparison between baseline and optimized scenarios for percentile 50

KPI	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday
Total Cost (€)	-53.78%	-62.21%	-57.98%	-57.04%	-59.95%	-58.88%	-59.98%
Distance travelled (km)	-13.29%	0.54%	-2.23%	-2.38%	-5.62%	-0.77%	-4.08%
Number of autonomous windows	-100.00%	-40.00%	25.00%	-25.00%	0.00%	50.00%	-50.00%
Pallets delivered with fleet	0.64%	1.50%	1.16%	1.00%	1.46%	1.21%	0.66%
Pallets delivered via subcontracted service	-100.00%	-100.00%	-98.51%	-100.00%	-100.00%	-98.65%	-100.00%
Vehicles required at peak	-20.00%	-9.52%	-9.76%	-1.20%	-9.20%	-8.14%	-10.39%
Maximum number of vehicles in preparation	-20.00%	2.38%	0.00%	6.98%	0.00%	-2.17%	-15.22%

Table B.2 - KPIs comparison between baseline and optimized scenarios for percentile 85

KPI	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday
Objective Function (€)	-53.59%	-57.47%	-59.70%	-58.62%	-69.39%	-65.34%	-65.91%
Distance Travelled (km)	-14.00%	-1.22%	1.18%	2.48%	1.65%	6.04%	3.25%
Autonomous Windows	-100.00%	-40.00%	75.00%	0.00%	-84.62%	75.00%	40.00%
Number of pallets delivered with fleet	0.65%	1.23%	1.93%	1.32%	2.55%	2.08%	1.07%
Number of pallets delivered via subcontracted service	-100.00%	-100.00%	-96.90%	-100.00%	-100.00%	-98.66%	-100.00%
Vehicles Required at peak	-15.79%	-12.36%	-4.82%	-7.53%	-11.96%	-5.00%	-6.10%
Maximum number of vehicles being prepared	-25.00%	-10.42%	-6.38%	-3.92%	-5.77%	-9.43%	-6.12%