Aalto University
School of Electrical
Engineering

Master's Programme in Mathematics and Operations Research

The combined second-echelon vehicle routing problem – A bi-objective approach

Patience Anipa

Master's Thesis
2025

Aalto University
School of Electrical Engineering

Author Patience Anipa

Title The combined second-echelon vehicle routing problem – A bi-objective approach

Degree programme Mathematics and Operations Research

Major Systems and Operations Research

Supervisor Philine Schiewe, Assistant Professor

Advisor Philine Schiewe, Assistant Professor

Date 31 July 2025 **Number of pages** 38 **Language** English

Abstract

Multi-objective optimization plays a crucial role in addressing complex decision-making problems where multiple, often conflicting objectives must be balanced simultaneously. In the realm of urban logistics, last-mile delivery represents a particularly challenging domain due to the need to optimize operational efficiency while minimizing environmental impact and congestion. Vehicle routing problems (VRPs) have long served as fundamental models to design effective delivery strategies, with recent advances incorporating multimodal and multi-echelon approaches to better reflect real-world logistics networks.

This thesis tackles the challenges of last-mile logistics in urban environments through the lens of multi-objective optimization, focusing on the Combined Second-Echelon Vehicle Routing Problem (CSERP)—a novel logistics framework that leverages existing public transportation infrastructure to reduce congestion and environmental footprint. In this setting, public transport vehicles (first-echelon), such as buses or trams, carry parcels to intermediate transfer points, where small, possibly autonomous second-echelon vehicles (e.g., drones, robots, or cargo bikes) complete the final leg of delivery.

To model the problem, a bi-objective optimization framework is developed to simultaneously minimize (1) the delay of the first-echelon vehicle and (2) the tour cost of the second-echelon vehicle. We consider a constrained setting in which a single second-echelon vehicle with unit capacity performs deliveries. Two scalarization techniques—weighted sum and ϵ-constraint methods—are applied to generate Pareto-optimal solutions, and we provide a theoretical discussion on the computational complexity and solvability of the resulting subproblems. The computational study used real-world data from the bus network in Göttingen, Germany, focusing on bus line 115 as a case study. The results obtained from the ϵ-constraint method supported the multi-objective optimization approach implementation for the combined second-echelon vehicle routing problem (CSERP). However, it was observed that the objectives are highly correlated. Therefore restricting the delay of the first-echelon vehicle further does not effect the cost of the second-echelon vehicle significantly, leading to only few and very similar points on the Pareto front.

Keywords Last-mile logistics, vehicle routing problem, two-echelon, scalarization, pareto, multicriteria optimization, bi-objective optimization

Contents

Abbreviations

MOO	Multi-objective Optimization
VRP	Vehicle Routing Problem
MEVRPs	Multi-Echelon Vehicle Routing Problems
2E-VRP	2-Echelon Vehicle Routing Problem

1 Introduction

The rapid growth of e-commerce has intensified competition among shippers to meet the increasingly personalized needs of customers. This surge in demand has also intensified the need for more sustainable and efficient solutions in city logistics, with last-mile delivery emerging as a critical area of focus (Kiba-Janiak et al. (2021)). Despite its importance, last-mile delivery is one of the least efficient segments of the supply chain, accounting for approximately 28% of total transportation costs (Ranieri et al. (2018)). The high costs and inefficiencies associated with this final stage of delivery have prompted companies and researchers to explore innovative strategies aimed at improving operational efficiency, optimizing delivery routes, and reducing environmental impact.

In an unpublished paper by Schiewe and Stinzendörfer (2024a), titled *The Combined Second-Echelon Routing Problem*, an innovative approach is presented to enhance last-mile logistics by utilizing spare capacity in public transport vehicles. Rather than relying solely on traditional delivery methods that deploy additional trucks and consequently increase traffic congestion in urban areas, their method integrates small, potentially autonomous vehicles—such as drones, robots, or cargo bikes—referred to as second-echelon vehicles, into existing public transport systems like trams, referred to as first-echelon vehicles. These first-echelon public transport vehicles transport packages to designated stops, where second-echelon vehicles take over the task of delivering the packages to their final destinations. The authors formulated the problem as a single-objective optimization problem, focusing on minimizing both the waiting times and the pure travel times of the second-echelon vehicles' delivery routes. Figure 1 illustrates the system with two second-echelon vehicles. The solid grey edges represent a line-bound vehicle route, while the dashed red and dotted blue edges depict the second-echelon routes serving eight customer locations with unit demand. Each second-echelon vehicle, with a capacity of two, serves two customers before returning to a handover station (represented by dark round nodes) for restocking.

The motivation for this thesis is drawn from the concepts presented in *The Combined Second-Echelon Routing Problem*. Their optimization framework aims to minimize the tour cost of the second-echelon vehicle, which implicitly minimizes the delay of the first-echelon vehicle. Thus, the optimization of the delay for the first-echelon vehicle is an unintended by-product. While effective in certain scenarios, this method lacks flexibility in managing the cost trade-off between the two vehicle types.

In many industrial settings, optimization problems are rarely defined by a single cost metric. Instead, multiple objectives must be considered. Even when cost is the primary focus, there are often different types of costs to evaluate, such as financial costs and time-related costs. Moreover, in vehicle routing problems, objectives usually extend beyond cost considerations to include factors such as fairness, punctuality, and customer satisfaction. This growing complexity has driven increased interest in multi-objective vehicle routing problems, where trade-offs between competing goals are systematically analyzed (Jozefowiez et al. (2008a)).

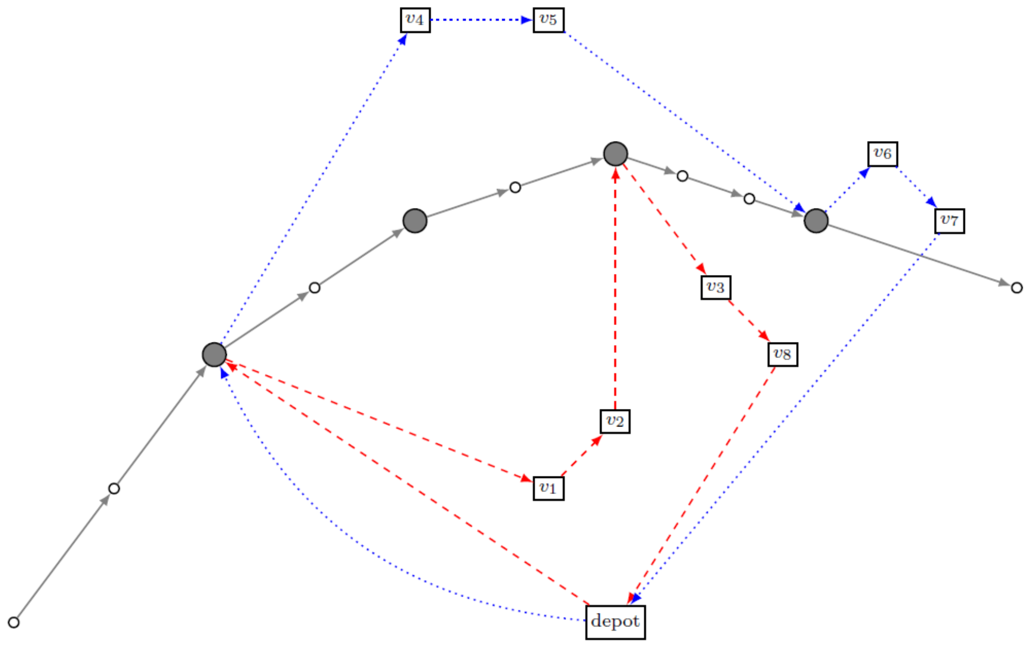

Figure 1: Tours of two second-echelon vehicles (dashed red edges and dotted blue edges, respectively) for a given public transport line (solid grey edges). Possible restocking stations are marked by dark circles while customer locations are marked by rectangles.

A key example of this complexity arises in the distribution of perishable goods and vaccines, where balancing cost and timely delivery is critical. In such cases, even if cost minimization is a priority, delays in the first-echelon vehicle can have a significantly greater negative impact on customer satisfaction and service quality. Therefore, a more holistic approach is needed—one that explicitly accounts for both cost and first-echelon vehicle delays, along with other operational factors, to accurately model and address real-world logistics challenges in these sensitive supply chains.

In this thesis, we extend the scope of the combined second-echelon routing problem by developing a bi-objective optimization model that minimizes the delay of the first-echelon vehicle and the tour costs of the second-echelon vehicle. The first-echelon vehicle typically operates on a fixed route and schedule, where delays can significantly impact overall system efficiency. The second-echelon vehicle cannot be transported, and its capacity is not limited to 1. By directly optimizing both objectives, this model allows for a more flexible and balanced trade-off between the performance of the two vehicles. The resulting solutions are expected to provide a set of Pareto optimal trade-offs, where neither objective can be improved without sacrificing the other, giving decision-makers the ability to select the most suitable solution based on specific operational priorities. To achieve this, we employ two scalarization techniques, the weighted sum scalarization method from the original model as a baseline and the

ϵ-constraint scalarization method. The complexity of the scalarization approaches will be analyzed for a simple case involving one second-echelon vehicle with a capacity of one. The results aim to advance the understanding of sustainable last-mile delivery solutions while contributing to the ongoing efforts to integrate public transport systems with modern logistics innovations.

1.1 Objectives

The main contributions of this thesis can be summarized as follows:

- We develop a bi-objective model of the combined second-echelon vehicle routing problem to explicitly address two key objectives: (1) minimize the delay of the first-echelon vehicle and (2) minimize the tour cost of the second-echelon vehicle,
- Considering the case of one second-echelon vehicle with a capacity of one, we implement the weighted sum and ϵ-constraint scalarization methods and conduct a detailed analysis on the theoretical properties i.e. whether they are polynomially solvable
- Utilizing real-world data, the Pareto front is computed by employing the ϵ-constraint scalarization method.

The remainder of the thesis is organized as follows: Chapter 2 presents the theoretical foundations of key concepts, including vehicle routing problems, multi-echelon vehicle routing problems with focus on the two-echelon variant and multi-objective optimization. Chapter 3 reviews the relevant literature, providing the foundational context for this research. Chapter 4 formally introduces the problem and its bi-objective formulations with one second-echelon vehicle. Chapter 5 explores various scalarization techniques for solving the bi-objective problem. Chapter 6 presents the experimental evaluation and analysis of the complexities of the scalarization techniques. Finally, Chapter 7 concludes the thesis by summarizing the findings and discussing potential directions for future research.

2 Theoretical Foundation

2.1 Vehicle routing problems

The origins of the Vehicle Routing Problem (VRP) trace back to the 1950s, though it stems from an even earlier problem—the Traveling Salesman Problem. The aim is to minimise costs, whether in terms of distance travelled, time or fuel consumed by designing optimal routes for a fleet of vehicles to serve a set of customers while adhering to specific constraints. It plays a crucial role in supply chain management by optimizing the physical delivery of goods and services. Various VRP variants exist, formulated based on factors such as the nature of the transported goods, service quality requirements, and the characteristics of both customers and vehicles. The VRP is classified as an NP-hard combinatorial optimization problem, meaning that solving it exactly for large real-world datasets within reasonable computational time is highly challenging. As a result, solution methods range from exact algorithms to heuristic and metaheuristic approaches, which provide near-optimal solutions more efficiently. Several VRP variations have been developed to incorporate real-world complexities, with the main ones being the Capacitated VRP (CVRP) which is the focus of this study, VRP with Time Windows (VRPTW), Pick-up and Delivery Vehicle Routing Problem (PDVRP) and the Multi-Depot Vehicle Routing Problem (MDVRP).

Also known as the classical VRP, the CVRP seeks to minimize cost by determining optimal delivery routes for vehicles with uniform characteristics under the constraints that each vehicle follows a single route which begins and ends at the depot, serves customers with known demands from a single central depot, visits each customer exactly once, and adheres to capacity constraints. The CVRP can be represented as a graph-theoretic problem, where a complete, undirected graph $G = (V, E)$ consists of a vertex set $V = \{0, \ldots, n\}$, with customer locations represented by $V_c = \{1, \ldots, n\}$ and the depot as vertex 0. Consider two connected nodes i and j; node i is connected to node j by an arc ij. If $c_{ij} \neq c_{ji}$, the problem is asymmetric (ACVRP); if $c_{ij} = c_{ji}$, it is symmetric (SCVRP). The cost c_{ij} is often assumed to be the Euclidean distance between customer locations, making the distance matrix symmetric and satisfying the triangle inequality:

$$c_{ik} + c_{kj} \leq c_{ij}, \quad \forall i, j, k \in V$$

The CVRP is known to be NP-hard. Exact algorithms used to solve the CVRP include branch-and-bound, branch-and-cut, and branch-and-price methods.

2.2 Multi-echelon vehicle routing problems

In practice, the distribution of the freight cannot be managed by direct shipping from the depot to the customers due to physical or legal restrictions, such as limited parking or vehicle weight limitations near customer locations. To ensure compliance with these constraints while preserving economies of scale, the distribution network is structured into multiple echelons, with goods transferred between various vehicle types or depots at different stages. In Multi-Echelon Vehicle Routing Problems (MEVRPs), freight delivery is optimized by rerouting and consolidating shipments between vehicles employed on either echelon through intermediate (satellite) locations (Guastaroba et al. (2016)). These problems typically involve capacity constraints on both vehicles and depots. The transportation network is structured into $k \geq 2$ levels:

- the first echelon, which connects the depots to the first-echelon intermediate locations
- $k - 2$ intermediate echelons interconnecting the intermediate locations
- the last echelon, where the freight is delivered from the intermediate locations to the customers.

When freight is consolidated from the depot to a satellite and then delivered from the satellite to the customer, it implicitly defines a two-echelon transportation system: the first level connects the depot to the satellites, while the second level connects the satellites to the customers. This variant of MEVRPs is known as the Two-Echelon Vehicle Routing Problem (2E-VRP). There are several variants of the 2E-VRP, which can be categorized into three main groups based on the nature of time dependence. These groups include basic variants with no time dependence, variants with time-dependent factors, and others that introduce further complexity, such as specific vehicle capacity constraints or multi-depot considerations. In the context of this study, we focus on the Two-Echelon Capacitated Vehicle Routing Problem (2E-CVRP), which operates as follows:

- Freight arrives at an external zone, the depot, where it is consolidated into the first-echelon vehicles, unless it is already carried in a fully-loaded first-echelon truck;
- Each first-echelon vehicle travels to a subset of satellites and then returns to the depot;
- At a satellite, freight is transferred from first-echelon vehicles to second-echelon vehicles;
- Each second-echelon vehicle performs a route to serve the designated customers and then travels to a satellite for its next cycle of operations. The second-echelon vehicles return to their departure satellite.

The problem is easily seen to be NP-Hard via a reduction to VRP, which is a special case of 2E-CVRP arising when just one satellite is considered.

3 Literature review

This chapter presents a comprehensive review of the relevant literature, building on the key topics and concepts introduced in Chapter 2. The review begins by examining foundational studies that establish the business case for this research, emphasizing the importance of vehicle routing problems (VRPs) and the role of multi-objective optimization (MOO) in addressing their complexities. Following this, the application of MOO in multi-echelon vehicle routing problems (ME-VRPs), particularly two-echelon vehicle routing problems is explored. Lastly, the chapter presents studies that apply MOO to problem settings closely aligned with the one investigated in this thesis.

3.1 Vehicle routing problems

Academic interest in the Vehicle Routing Problem (VRP), one of the most widely studied topics in Operations Research, along with its various adaptations, has steadily increased over the past several decades (Braekers et al. (2015)). The Truck Dispatching Problem, introduced by Dantzig and Ramser (1959), laid the foundation for algorithmic approaches to routing problems by formulating a mathematical model to optimize gasoline deliveries to service stations. A few years later, Clarke and Wright (1964) introduced a more efficient greedy heuristic that refined the original approach, transforming it into a widely applicable linear optimization model for logistics and transportation. These pioneering studies spurred extensive research, leading to the development of hundreds of models and algorithms aimed at finding both optimal and approximate solutions for various VRP variations.

However, modern VRP models have evolved significantly and now differ greatly from those introduced by Dantzig and Ramser (1959), and Clarke and Wright (1964). They incorporate real-world complexities such as time-dependent travel times influenced by traffic congestion, time windows for deliveries and pickups, and dynamically changing input data, such as fluctuating demand—all of which add substantial complexity to the problem.

3.2 Multi-objective optimization in vehicle routing problems

Multi-objective optimization in vehicle routing is not a recent development; research in this field dates back to the 1980s. One of the earliest studies by Park and Koelling (1986) addressed a standard vehicle routing problem (VRP) using a goal programming approach to minimize total distance, maximize fulfillment, and reduce goods deterioration. Over the following two decades, numerous studies tackled similar challenges, culminating in an extensive literature review by Jozefowiez et al. (2008b) that summarizes the key advancements in multi-objective vehicle routing problems.

Jozefowiez et al. highlight that traditional vehicle routing problems (VRPs) are typically formulated as single-objective optimization problems, primarily aiming to minimize costs or travel distance. However, real-world logistics often involve multiple conflicting objectives, such as balancing cost efficiency with service quality, fairness, and environmental impact. Their work situates these challenges within the framework of multi-objective optimization (MOO), where additional objectives can be incorporated to account for factors beyond cost minimization while preserving the core problem structure. This underscores the growing significance of MOO in VRPs due to the increasing complexity of practical applications.

A survey by Sandhya and Goel (2018) explores the role of multi-objective vehicle routing problems (MOVRPs) in three main ways. First, they extend classical VRPs by incorporating additional objectives while preserving the original problem structure, allowing researchers to explore new optimization goals without altering fundamental constraints. Second, MOVRPs generalize VRPs by transforming constraints into objective functions, shifting the focus from constraint satisfaction to multi-objective optimization. Lastly, they serve as models for real-world logistics challenges where decision-makers explicitly define multiple objectives to improve operational efficiency.

The new objectives introduced in the literature can be classified into tour-related, node/arc-related, and resource-related categories. Tour-related objectives typically focus on minimizing costs, reducing make-span, and balancing workloads. Node/arc-related objectives often involve optimizing time windows, improving customer satisfaction, and strengthening driver-customer relationships. Resource-related objectives primarily aim to minimize vehicle usage, providing both economic and environmental benefits, while also reducing goods damage. Solution approaches are generally divided into scalar methods such as weighted aggregation and goal programming, Pareto-based methods that leverage Pareto dominance, and non-scalar or non-Pareto methods, including lexicographic strategies and heuristics.

3.3 Optimization in multi-echelon vehicle routing problems

The multi-echelon vehicle routing problem represents a key extension of last-mile delivery challenges (Grangeon et al. (2008); Perboli et al. (2011)), aiming to optimize the movement of goods from depots to customers through multiple distribution stages. One of the most widely explored variants within this framework is the two-echelon vehicle routing problem (2E-VRP), which involves coordinating deliveries across two interconnected routing levels (Caggiani et al. (2015a)).

The 2E-VRP generally involves a two-stage distribution process: goods are first transported from a central depot to intermediate facilities, forming the first echelon. In the second echelon, deliveries are made from these intermediate facilities to the final customers. Although numerous studies have explored multi-objective approaches in the classic vehicle routing problems, research on the implementation of multi-objective optimization in the two-echelon vehicle routing problem (2E-VRP) remains relatively limited. Caggiani et al. (2015b) conducted a comprehensive review of the variants of the two-echelon vehicle routing problem (2E-VRP), while Cattaruzza et al. (2017) provided an overview of multi-level distribution systems in urban logistics. However, neither incorporates multi-objective optimization.

A growing area of interest within the two-echelon vehicle routing problem (2E-VRP) is the integration of electric vehicles, first introduced by Breunig et al. (2017) and further explored in subsequent studies (Breunig et al. (2019); Jie et al. (2019)). The first comprehensive review of two-echelon electric vehicle routing problems was conducted by Moradi et al. (2024a), providing an overview of the evolution of this research area. Although research output declined in subsequent years, this research area has continued to attract significant attention due to its potential for optimizing sustainable logistics, with vehicle-drone collaboration emerging as a notable variant. Nevertheless, despite these advancements, the application of multi-objective optimization in this domain remains largely unexplored.

Minimizing costs is the most widely used objective function in the literature and has been extensively studied in various works, including the previously mentioned research by Breunig et al. (2017), Breunig et al. (2019), and Jie et al. (2019), as well as Agárdi et al. (2019), Wang et al. (2019), Affi (2020), Wang and Zhou (2021), Zijlstra et al. (2021), and Wu and Zhang (2023). These authors mostly addressed objective functions such as minimizing traveled distances, fixed vehicle usage costs, fixed satellite utilization costs and energy/battery consumption costs, including battery swapping. While cost minimization is often optimized alongside other objectives, no study has explicitly addressed delivery tardiness or late deliveries within their models. Instead, these studies have incorporated time window-related constraints. Delivery tardiness, which can lead to penalties or reduced customer satisfaction, remains a critical yet underexplored challenge in two-echelon systems (Moradi et al. (2024b)).

A generalized variant of the 2E-VRP, similar to the model in this study, is presented by Schiewe and Stinzendörfer (2024b). A hybrid truck-cargo bike model is used, where the truck both delivers packages and serves as a mobile mini-depot. When the bike runs out of packages, both vehicles meet at a customer location to reload and continue simultaneous deliveries. While the model by Schiewe and Stinzendörfer (2024b) optimizes routes for both vehicles, the model in this study represents a special case where the truck's route and handover locations are predefined, and only the bike's tour is optimized.

Based on the literature, the two-echelon vehicle routing problem (2E-VRP) has been extensively studied; however, to the best of our knowledge, no multi-objective 2E-VRP model has been proposed that integrates public transportation in the first echelon and small, potentially self-driving vehicles in the second echelon, focuses on optimizing the second-echelon, while simultaneously optimizing the delay of the first-echelon vehicle and the tour cost of the second-echelon vehicle as distinct objectives. Furthermore, the application, comparison and complexities of various scalarization techniques of this specific bi-objective model has not been explored. Therefore, the primary contribution of this research is the development of a novel bi-objective optimization model aimed at minimizing both the delay in the first echelon and the tour cost in the second echelon. The secondary contribution is the application and comparison of different scalarization methods to solve the model, providing insights into the trade-offs between the objectives through the use of Pareto-optimal solutions.

4 Problem description and model formulation

The problem is a variant of the Capacitated Vehicle Routing Problem (CVRP), where for each tour, all customers have to be visited, the demands are known, all vehicles are identical and they all belong to the same depot. The objective is to minimize the total travel cost, which is the sum of each route's cost. Each route is a vehicle tour such that:

i. each tour starts and ends at the depot;

ii. each customer is visited once;

iii. the sum of the demands of the customers visited in a tour does not exceed the vehicle capacity.

As stated in the introduction, the combined second-echelon vehicle routing problem includes two echelons of transportation. The first-echelon vehicle is a public transportation vehicle known as a line-bound vehicle, which has fixed stations and follows a fixed schedule to make all deliveries. The nodes of the line-bound vehicle are denoted by the set $\{\tau_1, \dots, \tau_\ell\} := L$ and the corresponding scheduled departure times are $t(\tau_1) < \cdots < t(\tau_\ell)$. In defining the set of nodes of the first-echelon vehicle tour, we add a node for the depot, $DEP_f = \tau_0 = \tau_{\ell+1}$. The first-echelon vehicle tour is then given by $F = L \cup \{DEP_f\}$. The hand-over locations of the first echelon vehicle's tour (i.e. the locations where the second-echelon vehicle meets the first-echelon vehicle to pick up the packages), can be represented by the set $\{\kappa_1, \dots, \kappa_\ell\} \subset L$. These nodes are also known as the combined nodes. $F = (DEP_f = \tau_0, \tau_1, \dots, \tau_\ell, \tau_{\ell+1} = DEP_f)$ denotes the first-echelon vehicle tour.

Given that every pair of distinct vertices is connected by a pair of unique edges (one in each direction), we present the tour of the second-echelon vehicle as a complete digraph $G = (V, E)$ with $V = L \cup S \cup \{\text{DEP}\}$, where $S := \{v_{s1}, \dots, v_{sn}\}$ denotes the set of customer locations and $\text{DEP} = s_0 = s_{\ell+1}$ represents the depot. The edge costs $c(e)$, $e \in E$ and capacity $C \in \mathbb{N}_{>0}$ of the second-echelon vehicle are given, together with the demand of all customers $v \in S$. It is assumed that the cost function c satisfies the triangle inequality in G, where the shortest path cost between two nodes is stored as the corresponding edge weight. Moreover, we assume the demand of each customer $v \in S$ is equal to one, the second-echelon vehicle cannot be transported, and its capacity is limited to 1. The second-echelon vehicle tour is denoted by $S = (\text{DEP} = s_0, s_1, \dots, s_\ell, s_{\ell+1} = \text{DEP})$. The second-echelon vehicle has to meet up with the line-bound vehicle at the hand-over locations regularly in order to be reloaded with the required deliveries. Since it starts without any load, the first visited node after the depot has to be a hand-over location, therefore we assume $s_1 = \tau_1$. For tour S to be feasible, the tour has to cover the demand of all nodes, i.e., each node has to be served by the second-echelon vehicle. We assume that the both tours F and S are non-empty.

Problem 2. The *combined second-echelon routing problem (CSERP) using the bi-objective approach*, is to find a feasible tour $S = (\text{DEP} = s_0, s_1, \dots, s_\ell, s_{\ell+1} = \text{DEP})$ of the second-echelon vehicle such that the delay of the first-echelon vehicle and the generalized costs of the second-echelon vehicle tour are minimized.

4.1 Cost structure

In transport economics, generalized cost refers to the total cost of a journey, encompassing both monetary expenses (like fares, fuel, and tolls) and non-monetary costs (such as travel time or distance). The most important distinction we make is between independent costs and synchronized costs, similar to Schiewe and Stinzendörfer (2024b), which are used to represent different parts of the objectives.

Definition 3. (Independent costs [Schiewe and Stinzendörfer (2024b)]). The independent costs are the pure driving time or distance (i.e., a weighted sum of the distance covered) covered by the first- and second-echelon vehicles, respectively. For some objectives, these costs can be calculated individually and allocated to the relevant edges as $c(e)$. The route of the first-echelon vehicle has fixed stations and schedules, which means that the specific sequence of stops and edges in the first echelon's tour are predetermined and do not vary during optimization. When the first-echelon route is fixed, the objective function no longer needs to account for the optimization of these routes. The fixed independent costs for the first-echelon vehicle tour $F = (\tau_1, \ldots, \tau_l, \tau_{l+1} = DEP_\ell)$ can be formulated as:

$$c_f(F) := \sum_{i=1}^{l} c_f(\tau_i, \tau_{i+1})$$

As previously mentioned, a hand-over location must be the second-echelon vehicle's first visited node following the depot. The associated edge (DEP, s_1) is not included in the optimization problem and is fixed. Therefore, we define the costs of the second-echelon vehicle tour $S = (s_1, \ldots, s_l, s_{l+1} = \text{DEP})$ as follows, excluding the associated weight

$$c_s(S) := \sum_{i=1}^{l} c_s(s_i, s_{i+1})$$

Definition 4. (Synchronized costs [Schiewe and Stinzendörfer (2024b)]). The completion time of the second-echelon vehicle and the delay of the first-echelon vehicle, or the time both vehicles meet at the last combined node κ_ℓ, which is at least $t(\kappa_\ell)$, must be modeled using synchronized costs. This is due to the fact that we need to account for their waiting times at the hand-over locations. For the first-echelon vehicle, we extend the notation and define the synchronized costs for a path from node τ_i to node τ_j (where $i < j$) regarding the first-echelon vehicle's tour $F = (\tau_1, \ldots, \tau_n)$ is given by:

$$c_f(\tau_i, \tau_j, F) := \sum_{h=i}^{j-1} c_f(\tau_h, \tau_{h+1}),$$

Similarly the second echelon-vehicle costs of a path from node s_i to node s_j (where $i < j$) regarding tour $S = (s_1, \ldots, s_l, s_{l+1} = \text{DEP})$ by:

$$c_s(s_i, s_j, S) := \sum_{h=i}^{j-1} c_s(s_h, s_{h+1}),$$

It is necessary to explicitly simulate the coordination of both tours rather than treating the optimization of one or both of their durations as an independent objective. The synchronization of both tours must be ensured at each combined stop; that is, the vehicles must wait for each another because the duration to the last combined node must be long enough for both vehicles to visit all intermediate stops. The duration between two successive combined nodes κ_i and κ_{i+1} on the corresponding tour, is the cost of the slower vehicle (the vehicle with the higher summed costs) between the two nodes. To determine the costs of the tour, we need to sum up these durations between all successive combined nodes, as well as the summed costs (of the corresponding vehicle) between the last combined node κ_l and the depot.
We assume that both vehicles arrive to the first combined node κ_1 simultaneously and do not include the expenses of the depot's fixed outgoing edge, as we did in the independent case of the second-echelon vehicle. For the first- and second-echelon vehicle tours, respectively, this leads to the following definitions:

$$\tilde{c_f}(F) := \sum_{i=1}^{l-1} \max \left\{c_f(\kappa_i, \kappa_{i+1}, F), c_s(\kappa_i, \kappa_{i+1}, S)\right\} + c_f(\kappa_l, DEP_\ell, F) \tag{1}$$

$$\tilde{c_s}(S) := \sum_{i=1}^{l-1} \max \left\{c_f(\kappa_i, \kappa_{i+1}, F), c_s(\kappa_i, \kappa_{i+1}, S)\right\} + c_s(\kappa_l, \text{DEP}, S) \tag{2}$$

The first edge of the second-echelon vehicle tour is not part of the optimization problem.

4.2 MIP Formulation

In this section, the MIP formulation is presented with one first-echelon vehicle and one second-echelon vehicle based on independent and synchronized (time-based) costs. The objectives are given as follows:

$$\min \quad f_1(x) = \text{Delay of first-echelon vehicle}$$

$$\min \quad f_2(x) = \text{Tour costs of second-echelon vehicle}$$

where x represents the decision variables, including the vehicle routes.
We define binary variables $x_{(v,w)}$, which indicate whether the corresponding edge $(v, w) \in E$ is used by a second-echelon vehicle. The variables d_v represent the costs of the second-echelon tour starting at τ_0 or s_0 up to node $v \in V$, taking into account that both vehicles have to wait for each other at the combined nodes. More intuitively, d_v represents the time when the second-echelon vehicle reaches node v. Lastly, we define the variables ℓ_v for $v \in V$, which represent the number of goods delivered by

the second-echelon vehicle to node v and all preceding nodes, starting after the last reloading at a combined node. This accounts for the limited capacity C. The resulting MIP is shown in (3).
The first objective aims to minimize delays in the first-echelon vehicle's tour, ensuring precise synchronization with the second-echelon vehicle. The completion time of the second-echelon vehicle and the delay of the first-echelon vehicle, i.e., the time both vehicles meet at the last combined node κ_l, is at least $t(\kappa_l)$. The variable d_{κ_ℓ} represents the time when the second-echelon vehicle reaches the last hand-over location κ_ℓ. Therefore the delay of the first-echelon vehicle is formulated as the difference between the time when the second-echelon vehicle reaches the last hand-over location and the expected arrival time as shown in (3a). In the second objective, (3b), the goal is to determine a feasible tour for the second-echelon vehicle, $S = (\text{DEP} = s_0, s_1, \dots, s_l, s_{l+1} = \text{DEP})$, that minimizes the generalized costs of the tour. This is formulated as the weighted sum of the time-dependent costs of the second-echelon vehicle (i.e., d_{DEP}) represented in the first part of the objective function and the costs that are independent of waiting times (e.g., pure travel times), represented by the total costs of all used edges in the second part.

$$\min \quad d_{\kappa_\ell} - t(\kappa_\ell) \tag{3a}$$

$$\min \quad \omega_1 \quad d_{\text{DEP}} \quad + \quad \omega_2 \sum_{\substack{e \in E \\ e \neq (\text{DEP}, \kappa_1)}} x_e \cdot c(e) \tag{3b}$$

s.t.

$$x_{(\text{DEP},\kappa_1)} = 1 \tag{3c}$$

$$\sum_{\substack{w \in V \\ w \neq v}} x_{(v,w)} = \sum_{\substack{w \in V \\ w \neq v}} x_{(w,v)} \leq 1, \quad \forall v \in L \tag{3d}$$

$$\sum_{\substack{w \in V \\ w \neq v}} x_{(v,w)} = \sum_{\substack{w \in V \\ w \neq v}} x_{(w,v)} = 1, \quad \forall v \in S \tag{3e}$$

$$\ell_v \leq C, \qquad \forall v \in S \tag{3f}$$

$$\ell_v + 1 - \ell_w \leq (1 - x_{(v,w)}) \cdot (C + 1), \qquad \forall (v, w) \in E, w \in S \tag{3g}$$

$$d_{\kappa_i} + t(\kappa_{i+1}) \leq d_{\kappa_{i+1}} + t(\kappa_i), \qquad \forall i \in [\ell - 1] \tag{3h}$$

$$d_v + c(v, w) \leq d_w + (1 - x_{(v,w)}) \cdot M, \qquad \forall (v, w) \in E, v \neq \text{DEP} \tag{3i}$$

$$x_e \in \{0, 1\}, \qquad \forall e \in E \tag{3j}$$

$$d_v, \ell_v \geq 0, \qquad \forall v \in V \tag{3k}$$

Constraints (3c) denotes the first fixed outgoing edge of the depot of the second-echelon vehicle, while (3d)-(3e) are flow conservation constraints which ensure that each first-echelon vehicle node $v \in L$ and each second-echelon vehicle node, $v \in S$, has exactly one incoming and outgoing edge on the tour respectively. Constraints (3f)–(3g) ensure that the second-echelon vehicle visits a line-bound vehicle node for reloading (if necessary) without exceeding the capacity, while (3h)–(3i) ensure that both vehicles meet at the same time at a combined node and take the resulting waiting times into account. Here, (3h) and (3i) consider the traveling time of the line-bound and second-echelon vehicle, respectively. Moreover, the latter serve as subtour elimination constraints for the tour. The constant M has to be chosen large enough; in particular, it has to satisfy

$$d_v - d_w + \max_{(v,w)\in E}\{c(v,w)\} \leq \max_{v\in V}\{d_v\} + \max_{(v,w)\in E}\{c(v,w)\},$$

if the corresponding $x_{(v,w)}$ is equal to zero.

5 Methodology

This section presents the methodology for solving the bi-objective combined second-echelon routing problem (CSERP). To tackle the problem's bi-objective nature, we apply scalarization—a multi-objective optimization (MOO) technique that transforms the bi-objective problem into a single-objective one. Specifically, we apply two scalarization techniques: the weighted-sum method and the ϵ-constraint method. We consider the case of one second-echelon vehicle when the vehicle capacity is one and show that the resulting second-echelon scheduling problem is solvable in polynomial time. A comparative analysis of the scalarization approaches and solution methods is also provided.

5.1 Weighted Sum Scalarization

From the assumptions established in Section 4, we draw the following implications for the problem structure. Each customer node is visited exactly once, and the second-echelon vehicle is assumed to serve at least one customer in each segment between two combined nodes, as well as between the final combined node and the depot. This presupposes that the number of available customer nodes is sufficient to construct such segments; otherwise, any excess or unassigned nodes can be disregarded. Additionally, we assume that the sets of first- and second-echelon vehicle nodes are of equal size, i.e., $|F| = |S| = n$. If this condition is not initially satisfied, auxiliary second-echelon customer nodes with zero-cost assignments may be added to achieve balance. The order of the combined nodes in the second-echelon route S becomes fixed and known if the system does not allow intermediate deposition or storage of deliveries between successive combined nodes. In such a case, the second-echelon vehicle must visit the combined nodes in the exact order that matches the sequence of customer deliveries without any flexibility to reorder or delay hand-over operations. Lastly, the second-echelon vehicle is required to begin its tour at a combined node κ_1.

We now apply the scalarization technique known as the weighted sum method to reformulate the biobjective problem into a single-objective optimization problem [Marler and Arora (2004)]. In this approach, each objective function is assigned a non-negative scalar weight that encodes its relative importance in the trade-off. Given objective functions $f_1(x), f_2(x), \ldots, f_k(x)$, the scalarized objective is defined as

$$\min_{x \in X} \sum_{i=1}^{k} \lambda_i f_i(x),$$

where $\lambda_i \in \mathbb{R}_{>0}$ for all $i \in \{1, \ldots, k\}$, and the weights are normalized such that $\sum_{i=1}^{k} \lambda_i = 1$.

This normalization ensures that the scalarized function remains within the same objective space and facilitates the interpretation of each λ_i as the relative contribution of the i-th objective to the overall performance. Geometrically, the set of all possible weight vectors lies on the $(k-1)$-dimensional unit simplex in $\mathbb{R}^k$, which ensures convex combinations of objectives. Varying the weight vector $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_k)$ allows us to trace different parts of the Pareto front. In practical applications, a finite set of normalized weight vectors is typically sampled to approximate the full trade-off surface, particularly when exact preference information is unavailable or when the Pareto front needs to be estimated.

We define the weights $\lambda_i > 0, i \in \{1, ..., k\}$ and the weights are normalized such that $\sum_{i=1}^{k} \lambda_i = 1$. The weighted-sum objective of (3a) and (3b) is given by:

$$\min \quad \lambda_1 \cdot \left(d_{\kappa_\ell} - t(\kappa_\ell)\right) \quad + \quad \lambda_2 \cdot \left[\omega_1 \cdot d_{\text{DEP}} + \omega_2 \sum_{\substack{e \in E \\ e \neq (\text{DEP}, \kappa_1)}} x_e \cdot c(e)\right] \tag{4}$$

where $\lambda_2 = (1 - \lambda_1)$

For simplicity, we re-parameterize the weights as follows: $\theta_1 = \lambda_1$, $\theta_2 = \lambda_2\omega_1$, $\theta_3 = \lambda_2\omega_2$ and $\sum_{i=1}^{k} \theta_i = 1$. This gives the equivalent form:

$$\min \quad \theta_1 \cdot \left(d_{\kappa_\ell} - t(\kappa_\ell)\right) \quad + \quad \theta_2 \cdot d_{\text{DEP}} \quad + \quad \theta_3 \cdot \sum_{\substack{e \in E \\ e \neq (\text{DEP}, \kappa_1)}} x_e \cdot c(e) \tag{5}$$

Theorem 5. *If $C = 1$ holds, we can solve CSEVRP by computing a minimum-weight perfect matching, which can be done in polynomial time [Cook and Rohe (1999)].*

Proof. The fixed ordering of combined nodes in the second-echelon tour S imposes a structured relationship between hand-over locations and customer visits. In particular, for each segment of the tour indexed by $i \in [\ell]$, a customer node v_{s_i} must be associated with a preceding combined node κ_i. This sequence of associations can be abstracted as a pairing problem between two disjoint sets: the set of combined nodes and the set of customer nodes.

To formalize this, we construct a complete bipartite graph $G_B = (V_B, E_B)$, where the vertex set is given by $V_B = \{\kappa_1, \ldots, \kappa_\ell\} \cup \{v_{s_1}, \ldots, v_{s_\ell}\}$. A feasible routing plan then corresponds to a perfect matching M in G_B, such that its cardinality is equal to $|\{\kappa_1, \ldots, \kappa_\ell\}| = |\{v_{s_1}, \ldots, v_{s_\ell}|$ (i.e. there are exactly the same number of combined nodes as customer nodes).

The cost associated with each edge $(\kappa_i, u) \in E_B$ reflects the travel cost incurred when the second-echelon vehicle departs from combined node κ_i, visits customer node u, and proceeds toward the next combined node κ_{i+1} (or returns to the depot if $i = \ell$). These edge weights capture the second-echelon routing costs defined by the objective and are used to determine a minimum-cost matching that aligns with the fixed node sequence in S. For the tour $S = (s_1 = \kappa_1, s_2 = v_{s_1}, \ldots, s_{2\ell} = v_{s_\ell}, s_{2\ell+1} = \text{DEP})$, the weighted costs are:

$$\theta_1 \cdot \left(d_{\kappa_\ell} - t(\kappa_\ell)\right) + \theta_2 \cdot d_{\text{DEP}} + \theta_3 \cdot \sum_{i=1}^{2\ell} c(s_i, s_{i+1})$$

$$= (\theta_1 + \theta_2) \underbrace{\sum_{i=1}^{\ell-1} \max\left\{t(\kappa_{i+1}) - t(\kappa_i),\ c(\kappa_i, \mathrm{v}_{si}) + c(\mathrm{v}_{si}, \kappa_{i+1})\right\}}_{=d_{\kappa\ell}} - \theta_1 t(\kappa_\ell)$$

$$+\theta_2 \underbrace{\left(c(\kappa_\ell, \mathrm{v}_{s\ell}) + c(\mathrm{v}_{s\ell}, \text{DEP})\right)}_{=d_{\text{DEP}} - d_{\kappa\ell}}$$

$$+\theta_3 \sum_{i=1}^{\ell} \left(c(\kappa_i, \mathrm{v}_{si}) + c(\mathrm{v}_{si}, \kappa_{i+1})\right)$$

To minimize (5), we define the edge weights $w(\kappa_i, u)$ as:

$$(\theta_1 + \theta_2) \max\{\underbrace{t(\kappa_{i+1}) - t(\kappa_i),\ c(\kappa_i, u) + c(u, \kappa_{i+1})}_{\text{fixed}}\} + \theta_3 \underbrace{\left(c(\kappa_i, u) + c(u, \kappa_{i+1})\right)}_{\text{fixed}}, \quad i \in [\ell - 1] \tag{6}$$

$$(\theta_2 + \theta_3)\left(c(\kappa_\ell, u) + c(u, \text{DEP})\right) - \theta_1 t(\kappa_\ell), \quad i = \ell \tag{7}$$

Then calculate a minimum-weight perfect matching in G_B whose weight represents the total weighted costs of the second-echelon vehicle (3b).

Remark. *Since the term $\theta_1 \left(d_{\kappa_\ell} - t(\kappa_\ell)\right)$ in the weighted-sum formulation in (5) accounts for deviations from the scheduled timeline, the resulting negative term, $-\theta_1 t(\kappa_\ell)$, in the edge weight in (7) functions as a penalization mechanism, allowing the trade off between the delay of the first echelon vehicle and the generalized costs of the second echelon vehicle in the CSERP.*

We consider a basic example to demonstrate the use of this approach. We use two combined nodes κ_1 and κ_2 and two customer nodes v_1 and v_2 to ensure a perfect matching. The second-echelon vehicle cost of traveling between the nodes, $c(\kappa_i, v_{si})$ in time unit of minutes, and scheduled departure times of the line-bound vehicle from the combined nodes, $t(\kappa_i)$, are given as follows:

- $c(\kappa_1, v_1) = 10$
- $c(\kappa_1, v_2) = 15$
- $c(\kappa_2, v_1) = 12$
- $c(\kappa_2, v_2) = 8$

- $c(v_1, DEP) = 3$
- $c(v_2, DEP) = 18$
- $t(\kappa_1) = 2, t(\kappa_2) = 8$

The problem is visualized in the bipartite graph in Figure 2.

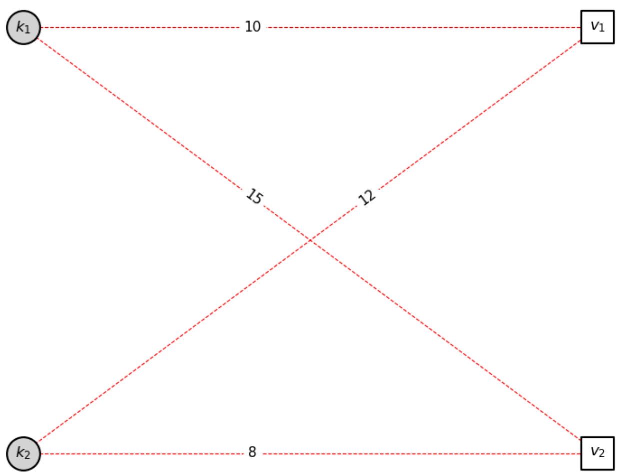

Figure 2: A bipartite graph showing tours of the second-echelon vehicle in dashed red edges with respective distance costs. Possible restocking stations are marked by dark circles while customer locations are marked by squares.

We set the weights for the scalarized objective (5) as:

$$\theta_1 = 0.7, \quad \theta_2 = 0.2, \quad \theta_3 = 0.1$$

The edge weights from each combined node to a customer node, $w(\kappa_i, u)$, are computed as follows using the formulas derived in (6) and (7):

For κ_1,

$$\begin{aligned} w(\kappa_1, v_1) &= (\theta_1+\theta_2) \max\{t(\kappa_2) - t(\kappa_1),\ c(\kappa_1, v_1) + c(v_1, \kappa_2)\} + \theta_3\left(c(\kappa_1, v_1) + c(v_1, \kappa_2)\right) \\ &= (0.7 + 0.2) \times \max\{6,\ 22\} + (0.1 \times 22) \\ &= (0.9 \times 22) + (0.1 \times 22) \\ &= 22 \end{aligned}$$

$$\begin{aligned} w(\kappa_1, v_2) &= (\theta_1+\theta_2) \max\{t(\kappa_2) - t(\kappa_1),\ c(\kappa_1, v_2) + c(v_2, \kappa_2)\} + \theta_3\left(c(\kappa_1, v_2) + c(v_2, \kappa_2)\right) \\ &= (0.7 + 0.2) \times \max\{6,\ 23\} + (0.1 \times 23) \end{aligned}$$

$$= (0.9 \times 23) + (0.1 \times 23)$$
$$= 23$$

For κ_2,

$$w(\kappa_2, v_1) = (\theta_2 + \theta_3)\,(c(\kappa_2, v_1) + c(v_1, \text{DEP})) - \theta_1 t(\kappa_2)$$
$$= (0.2 + 0.1) \times (15) - (0.7 \times 8)$$
$$= -1.1$$

$$w(\kappa_2, v_2) = (\theta_2 + \theta_3)\,(c(\kappa_2, v_2) + c(v_2, \text{DEP})) - \theta_1 t(\kappa_2)$$
$$= ((0.2 + 0.1) \times 26) - (0.7 \times 8)$$
$$= 2.2$$

The edge weights are shown in the following matching graph in Figure 3.

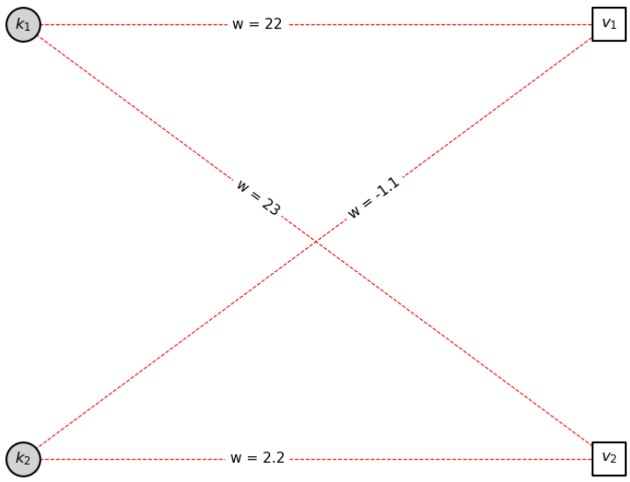

Figure 3: The bipartite graph is updated to reflect the computed edge weights $w(\kappa_i, u)$.

The objective is to identify a perfect matching that minimizes the total cost, i.e., the sum of edge weights. Consider the following feasible matchings and their associated total weights:

$$\text{Matching 1: } (\kappa_1, v_1), (\kappa_2, v_2) \quad \Rightarrow \quad \text{Total cost} = 22 + 2.2 = 24.2$$
$$\text{Matching 2: } (\kappa_1, v_2), (\kappa_2, v_1) \quad \Rightarrow \quad \text{Total cost} = 23 - 1.1 = 21.9$$

Since Matching 2 yields the minimum total edge weight, it is selected as the optimal solution. This matching represents a feasible assignment for the second-echelon vehicle that jointly minimizes its routing cost while satisfying the delay constraints imposed on the first-echelon vehicle.

5.2 ϵ-Constraint Scalarization

We consider an alternative scalarization approach known as the ε-constraint method, which minimizes one objective while treating the remaining objectives as inequality constraints [Haimes et al. (1971)]. Specifically, we minimize the cost associated with the second-echelon vehicle while imposing an upper bound $\varepsilon \in \mathbb{R}$ on the delay incurred by the first-echelon vehicle at each combined node κ_i. The delay constraint (8b) restricts feasible solutions to those where the delay at every handover location does not exceed the specified threshold ε. By systematically varying ε, we generate the Pareto frontier of trade-offs between cost and delay. In time-critical logistics systems, some delay may be unavoidable but must be tightly controlled to ensure effective subsequent operations. This reflects a common practical distinction: we are willing to pay more to stay within service constraints, but not to reduce delay beyond what is operationally necessary. This rationale also aligns with the structure of the ε-constraint method, which is particularly effective when one objective naturally acts as an operational limit [Mavrotas (2009)]. The resulting ε-constraint scalarization formulation of the CSEVRP is as follows:

$$\min \quad \omega_1 \cdot d_{\mathrm{DEP}} + \omega_2 \sum_{\substack{e \in E \\ e \neq (\mathrm{DEP}, \kappa_1)}} x_e \cdot c(e) \tag{8a}$$

$$\text{s.t.} \quad d_{\kappa_\ell} - t(\kappa_\ell) \leq \varepsilon \tag{8b}$$

$$x_{(\mathrm{DEP}, \kappa_1)} = 1 \tag{8c}$$

$$\sum_{\substack{w \in V \\ w \neq v}} x_{(v,w)} = \sum_{\substack{w \in V \\ w \neq v}} x_{(w,v)} \leq 1, \quad \forall v \in L \tag{8d}$$

$$\sum_{\substack{w \in V \\ w \neq v}} x_{(v,w)} = \sum_{\substack{w \in V \\ w \neq v}} x_{(w,v)} = 1, \quad \forall v \in S \tag{8e}$$

$$\ell_v \leq C, \quad \forall v \in S \tag{8f}$$

$$\ell_v + 1 - \ell_w \leq (1 - x_{(v,w)}) \cdot (C + 1), \quad \forall (v, w) \in E, w \in S \tag{8g}$$

$$d_{\kappa_i} + t(\kappa_{i+1}) \leq d_{\kappa_{i+1}} + t(\kappa_i), \quad \forall i \in [\ell - 1] \tag{8h}$$

$$d_v + c(v, w) \leq d_w + (1 - x_{(v,w)}) \cdot M, \quad \forall (v, w) \in E, v \neq \mathrm{DEP} \tag{8i}$$

$$x_e \in \{0, 1\}, \quad \forall e \in E \tag{8j}$$

$$d_v, \ell_v \geq 0, \quad \forall v \in V \tag{8k}$$

We use the problem definition from Theorem 5 and derive the edge weights for the objective function (8a) for the minimum weight perfect matching algorithm.

Consider the tours $F = (\mathrm{DEP}_f = \tau_0, \tau_1, \ldots, \tau_\ell, \tau_{\ell+1} = \mathrm{DEP}_f)$ and $S = (s_1 = \kappa_1, s_2 = v_{s_1}, \ldots, s_{2\ell} = v_{s_\ell}, s_{2\ell+1} = \mathrm{DEP})$. The objective (8a), representing the cost associated with the second-echelon vehicle, can be expressed as:

$$\omega_1 d_{\mathrm{DEP}} + \omega_2 \sum_{i=1}^{2\ell} c(s_i, s_{i+1})$$

$$= \omega_1 (\sum_{i=1}^{\ell-1} \max \{t(\kappa_{i+1}) - t(\kappa_i),\ c(\kappa_i, \mathrm{v}_{si}) + c(\mathrm{v}_{si}, \kappa_{i+1})\}$$
$$+ (c(\kappa_\ell, \mathrm{v}_{s_\ell}) + c(\mathrm{v}_{s_\ell}, \mathrm{DEP})))$$
$$+ \omega_2 \sum_{i=1}^{\ell} \left(c(\kappa_i, v_{s_i}) + c(v_{s_i}, \kappa_{i+1}) \right)$$

To minimize the objective (8a), we define the edge weights between combined nodes κ_i and intermediate customer nodes u as:

$$w(\kappa_i, u) := \begin{cases} \omega_1 \max \{ \underbrace{t(\kappa_{i+1}) - t(\kappa_i),\ c(\kappa_i, u) + c(u, \kappa_{i+1})}_{\text{fixed}} \} + \omega_2 \underbrace{(c(\kappa_i, u) + c(u, \kappa_{i+1}))}_{\text{fixed}}, & i \in \{\ell - 1\} \\ (\omega_1 + \omega_2)\,(c(\kappa_\ell, u) + c(u, \mathrm{DEP}))\,, & i = \ell \end{cases} \tag{9}$$

A minimum-weight perfect matching is then computed on the auxiliary bipartite graph G_B with edge weights defined in (9). The weight of this matching directly corresponds to the total weighted cost described in (8a).

5.2.1 Delay bounds $\varepsilon = \infty$ and $\varepsilon = 0$

We now analyze how the minimum-weight perfect matching problem incorporates the epsilon constraint (8b). Since the feasible region is determined by the choice of the vector ε, which bounds the allowable trade-off between the two objectives, we focus on the two extreme bounds (anchor points) on the Pareto front:

Case 1: $\varepsilon = \infty$

When ε is set to an arbitrarily large value, the delay constraint becomes non-binding. Specifically, since the constraint $d_{\kappa_\ell} - t(\kappa_\ell) \leq \varepsilon$ is always satisfied for any feasible matching when $\varepsilon = \infty$, it imposes no restriction on the solution space. Consequently, the delay term is effectively ignored, and the optimization problem reduces to minimizing the second-echelon vehicle cost :

$$
\begin{aligned}
\min \quad & \omega_1 \cdot d_{\text{DEP}} + \omega_2 \sum_{\substack{e \in E \\ e \neq (\text{DEP},\kappa_1)}} x_e \cdot c(e) && \text{(10a)} \\
\text{s.t.} \quad & x_{(\text{DEP},\kappa_1)} = 1 && \text{(10b)} \\
& \sum_{\substack{w \in V \\ w \neq v}} x_{(v,w)} = \sum_{\substack{w \in V \\ w \neq v}} x_{(w,v)} \leq 1, \quad \forall v \in L && \text{(10c)} \\
& \sum_{\substack{w \in V \\ w \neq v}} x_{(v,w)} = \sum_{\substack{w \in V \\ w \neq v}} x_{(w,v)} = 1, \quad \forall v \in S && \text{(10d)} \\
& \ell_v \leq C, \quad \forall v \in S && \text{(10e)} \\
& \ell_v + 1 - \ell_w \leq (1 - x_{(v,w)}) \cdot (C + 1), \quad \forall (v,w) \in E, w \in S && \text{(10f)} \\
& d_{\kappa_i} + t(\kappa_{i+1}) \leq d_{\kappa_{i+1}} + t(\kappa_i), \quad \forall i \in [\ell - 1] && \text{(10g)} \\
& d_v + c(v,w) \leq d_w + (1 - x_{(v,w)}) \cdot M, \quad \forall (v,w) \in E, v \neq \text{DEP} && \text{(10h)} \\
& x_e \in \{0,1\}, \quad \forall e \in E && \text{(10i)} \\
& d_v, \ell_v \geq 0, \quad \forall v \in V && \text{(10j)}
\end{aligned}
$$

This corresponds to the classical minimum-weight perfect matching problem, in which the second-echelon vehicle's cost is minimized with no consideration of the delay of the first-echelon vehicle.

Case 2: $\varepsilon = 0$

To enforce no delay for the first-echelon vehicle at all combined nodes κ_i, the vector ε may be set to zero, i.e., $\varepsilon = 0 \in \mathbb{R}$. In this case, only edges satisfying zero delay can be included in the matching. The edges (κ_i, u) in G_B for which the time taken by the first-echelon vehicle to travel from a combined node to the next combined node exceeds the total cost of a second-echelon vehicle tour from the same combined node to the other are removed from the tour. The restricted edge set is defined as:

$$E_0 = \{(\kappa_i, u) : t(\kappa_{i+1}) - t(\kappa_i) \geq c(\kappa_i, u) + c(u, \kappa_{i+1})\}$$

and the corresponding subgraph $G_0 = (V, E_0)$. The problem becomes:

$$
\begin{aligned}
\min \quad & \omega_1 \cdot d_{\text{DEP}} + \omega_2 \sum_{\substack{e \in E_0 \\ e \neq (\text{DEP},\kappa_1)}} x_e \cdot c(e) && \text{(11a)} \\
\text{s.t.} \quad & d_{\kappa_\ell} - t(\kappa_\ell) = 0 && \text{(11b)} \\
& x_{(\text{DEP},\kappa_1)} = 1 && \text{(11c)} \\
& \sum_{\substack{w \in V \\ w \neq v}} x_{(v,w)} = \sum_{\substack{w \in V \\ w \neq v}} x_{(w,v)} \leq 1, \quad \forall v \in L && \text{(11d)} \\
& \sum_{\substack{w \in V \\ w \neq v}} x_{(v,w)} = \sum_{\substack{w \in V \\ w \neq v}} x_{(w,v)} = 1, \quad \forall v \in S && \text{(11e)} \\
& \ell_v \leq C, \quad \forall v \in S && \text{(11f)} \\
& \ell_v + 1 - \ell_w \leq (1 - x_{(v,w)}) \cdot (C+1), \quad \forall (v,w) \in E, w \in S && \text{(11g)} \\
& d_{\kappa_i} + t(\kappa_{i+1}) \leq d_{\kappa_{i+1}} + t(\kappa_i), \quad \forall i \in [\ell - 1] && \text{(11h)} \\
& d_v + c(v,w) \leq d_w + (1 - x_{(v,w)}) \cdot M, \quad \forall (v,w) \in E, v \neq \text{DEP} && \text{(11i)} \\
& x_e \in \{0,1\}, \quad \forall e \in E && \text{(11j)} \\
& d_v, \ell_v \geq 0, \quad \forall v \in V && \text{(11k)}
\end{aligned}
$$

The zero delay matchings (if they exist) form a limited subset of all possible matchings, as the feasible region shrinks, often consequently leading to an empty feasible set. If no perfect matching exists in G_0, then the problem becomes infeasible under the bound $\varepsilon = 0$. This infeasibility indicates that the first-echelon vehicle cannot achieve zero delay even under an ideal matching, thereby establishing a fundamental lower bound on the achievable delay. The solution corresponding to $\varepsilon = 0$ represents the best-case delay scenario for the first-echelon vehicle, but is likely suboptimal for the second-echelon vehicle due to the constrained search space.

5.2.2 MIP formulation of the matching problem

We use the computed edge weights $w(\kappa_i, u)$ in (9) and define the binary variable $y(\kappa_i, u)$ which indicate whether a customer node u is visited between κ_i and κ_{i+1}. We also define the variable $d(\kappa_i, u)$ which represents the delay of the first echelon vehicle when a customer node u is visited between κ_i and κ_{i+1}.

$$
d(\kappa_i, u) = \begin{cases} max\big\{(c(\kappa_i, u) + c(u, \kappa_{i+1})) - (t(\kappa_{i+1}) - t(\kappa_i)), 0\big\} & i \in \{\ell - 1\} \\ 0, & i \in \{\ell\} \end{cases} \tag{12}
$$

The MIP formulation of the matching problem is given as follows:

$$\min \quad \sum_{(\kappa_i,u)\in E} w(\kappa_i,u)\cdot y(\kappa_i,u) \tag{13a}$$

$$\text{s.t.} \quad \sum_{(\kappa_i,u)\in E} d(\kappa_i,u)\cdot y(\kappa_i,u) \leq \varepsilon \tag{13b}$$

$$\sum_{\substack{\kappa_i\in J:\\(\kappa_i,u)\in E}} y(\kappa_i,u) = 1, \quad \forall u \in V \tag{13c}$$

$$\sum_{\substack{u\in V:\\(\kappa_i,u)\in E}} y(\kappa_i,u) = 1, \quad \forall \kappa_i \in J \tag{13d}$$

$$y(\kappa_i,u) \in \{0,1\}, \quad \forall (\kappa_i,u) \in E \tag{13e}$$

In the case where $\varepsilon = 0$, the matching is optimized over the restricted edge set E_0.

5.3 Algorithm

Based on the theoretical results in sections 5.1 and 5.2, we suggest the following algorithm for finding non-dominated solutions to the bi-objective problem combined second-echelon routing problem (CSERP) for one second echelon vehicle with capacity one $C = 1$. The problem is solved using one of the presented scalarization methods, the ϵ-constraint scalarization method, which guarantees that all non-dominated solutions can be found, thus this will be investigated. For the given line-bound vehicle tour F, combined nodes $(\kappa_1,\dots,\kappa_\ell)$ and set of customers $\{v_{s1},\dots,v_{s\ell}\}$, we define the complete bipartite graph $G_B = (V_B, E_B)$ with $V_B = \{\kappa_i : i = 1,\dots,\ell\} \cup \{v_{si} : i = 1,\dots,\ell\}$. As stated in the proof of Theorem 5, any feasible solution is an assignment of a customer v_{si} to the preceding combined node κ_i for all $i \in [\ell]$. Consequently, the edge weights $w(\kappa_i,u)$ in G_B for $\kappa_i \in \{\kappa_i : i = 1,\dots,\ell\}$ and $u \in \{v_{si} : i = 1,\dots,\ell\}$ correspond to the weighted second-echelon vehicle costs (driving from node κ_i via u to κ_{i+1}) with respect to the objective function. Those weights are defined as described in (9). Algorithm 1 shows the complete procedure.

Algorithm 1 ϵ-constraint scalarization

1: Define $G_B = (V_B, E_B)$, weights $w(\kappa_i,u)$ for $\kappa_i \in \{\kappa_i : i = 1,\dots,\ell\}, u \in \{v_{si} : i = 1,\dots,\ell\}$, $\varepsilon = \{\varepsilon_1,\dots,\varepsilon_m\}$
2: Initialize $P \leftarrow \emptyset$
3: **for** $\varepsilon_i \in \varepsilon$ **do**
4: Compute ϵ-constraint scalarization on G_B (using MIP in section 5.2.2) with optimal solution λ_ε^*
5: Compute $y_{\varepsilon_1} \leftarrow f_1(\lambda_\varepsilon^*)$ and $y_{\varepsilon_2} \leftarrow f_2(\lambda_\varepsilon^*)$
6: $P \leftarrow P \cup \{(y_{\varepsilon_1}, y_{\varepsilon_2})\}$
7: **end for**
8: return

6 Experimental evaluation

This section presents the implementation of the algorithm from Section 5.3 on real-world data, with the aim of evaluating the resulting Pareto front and demonstrating the practical effectiveness of the proposed ϵ-constraint model. For the computational study, we restrict ourselves to the objective function with weights $\omega_1 = 0.2$ and $\omega_2 = 0.8$ as this minimizes the makespan of the second-echelon vehicle (d_{DEP}).

The algorithm has been implemented in `Python`, and tests were performed on a 64-bit system with an Intel® Core™ i5-6300U CPU at 2.40 GHz and 8 GB of memory.

6.1 Data Description

This computational analysis utilizes the *goevb* dataset from the LinTim scientific software toolbox [Schiewe et al. (2024)], which models the public bus network in Göttingen, Germany. Given the lack of existing benchmarks tailored to our model, this dataset offers the most appropriate foundation. Bus line 115 was chosen for its strong resemblance to the operational scenario under consideration. As illustrated in Figure 4, the route originates in the northwestern outskirts of the city, proceeds through the central district, loops around in the eastern neighborhoods, and then retraces its path back. The violet marker in the upper left of the map indicates the departure point (depot) of the line-based vehicle. Meanwhile, the depot for the second-echelon fleet is strategically positioned in the first suburban area, adjacent to a supermarket with enough space for vehicle dispatch.

A total of 32 station locations are marked with small blue dots, while customer nodes—selected randomly—are shown in red. Since the bus route involves a turnaround at its final stop, the full path includes 61 service points, representing the complete set of nodes along the outbound and return segments.

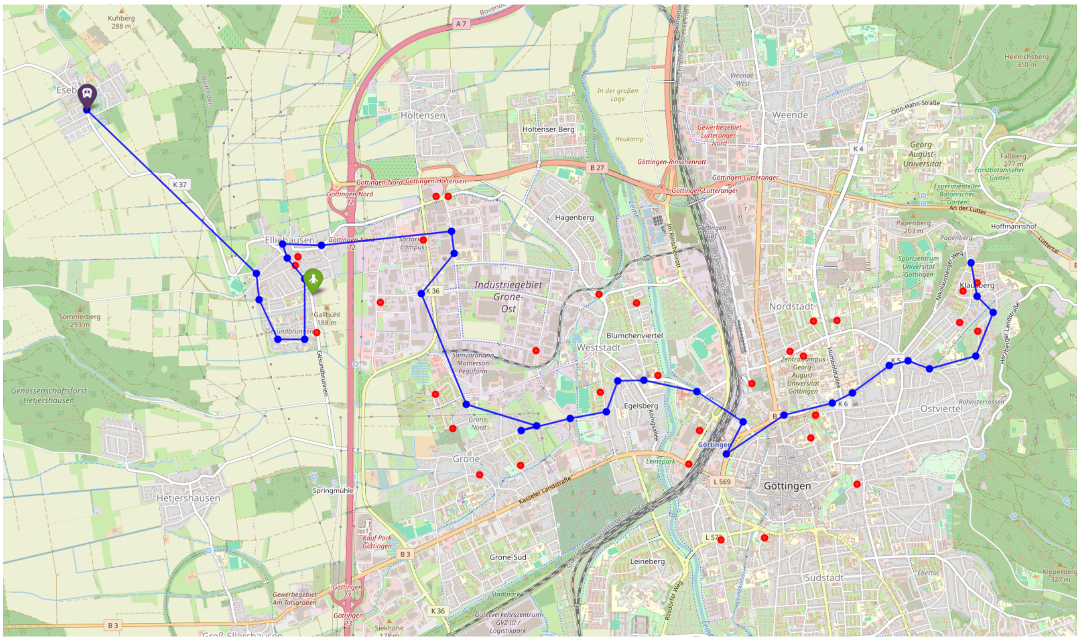

Figure 4: Visualization of the Göttingen area showing the path of bus line 115, including its 32 stations (indicated by small blue markers), the origin of the line-bound vehicle (highlighted with a large violet marker in the upper left), the second-echelon depot location (shown as a large green marker positioned in the early section of the route), and 32 randomly selected subset of customer locations represented by small red dots.

In contrast to much of the existing literature, which typically relies on Euclidean distances—particularly in drone-based routing—we employ the actual path network to compute distances. This approach enables the inclusion of various second-echelon vehicle types beyond drones. Although delivery drones can reach speeds of up to 128 km/h, our model incorporates additional factors such as takeoff, arrival, and package drop-off times. As a result, we assume an average effective drone speed of 60 km/h.

6.2 Results

In this section the results obtained from implementing the ε-constraint method is presented. In order to analyze the trade-off between the two objectives; minimizing the delay of the first-echelon vehicle and the generalized tour cost of the second-echelon vehicle, we compute the Pareto front. The formulated a Mixed-Integer Programming (MIP) model in section 5.2.2 is implemented to generate the set of Pareto points. The model includes 65 constraints, 1,024 binary decision variables, and 2,961 non-zero elements, reflecting the complexity of the distribution network and the granularity of the decision space.

The optimization model was implemented using *Python 3.8.12* and its optimization library *Python MIP*. Optimizations were solved successfully solved using the CBC solver. Despite the model's size, the solver handled it efficiently. Conflict graph and symmetry detection techniques were applied to simplify the model structure, but no significant reductions were obtained, indicating the formulation did not contain redundant or easily reducible symmetries. Importantly, the feasibility pump heuristic of the solver succeeded in finding a feasible solution with an objective value of 72.8 within just 0.02 seconds. This performance, notably achieved without requiring branch-and-bound iterations, shows both the tractability of the model and that the heuristic was effective for this instance.

As mentioned in the previous section, in this approach, the the tour cost of the second-echelon vehicle is minimized, while the delay of the first-echelon vehicle is bounded by an ε threshold. The ε values were systematically varied over a range to explore the trade-off surface between these two objectives. We used a fixed stepwise increment for ε values, ranging from 32 to 76 minutes. The tour cost was minimized under each of these delay constraints, and the resulting objective pairs were plotted to form the Pareto front. Despite the broad sweep of allowable delay bounds, the model returned only a few distinct non-dominated solutions. These are depicted in the Pareto front graph in figure 5, where the horizontal axis represents the delay of the first-echelon vehicle $f_1(x)$ and the vertical axis represents the cost of the second-echelon tour $f_2(x)$.

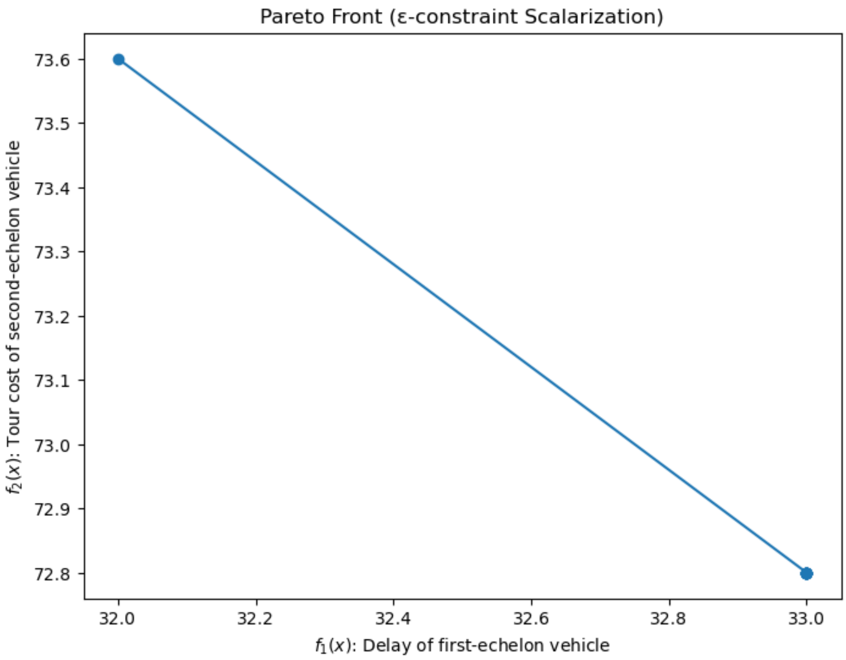

Figure 5: Each point represents a non-dominated solution: improving one objective necessitates degrading the other. For example, minimizing the delay to its lowest value (around 32 minutes) results in a cost of around 73.6, while increasing the delay to about 33 minutes reduces the cost to approximately 72.8. These changes, although numerically small, are significant in operational terms, highlighting the potential impact of even minute planning decisions.

As shown in the figure above, the resulting Pareto front is composed of only two closely clustered points. This suggests that most delay constraints (i.e., ε values) yielded identical or nearly identical tour costs, with variation occurring only in a narrow subregion around $\varepsilon \approx 32$–33 minutes.

This implies a high degree of correlation between the two objectives. That is, minimizing tour cost tends to simultaneously yield minimal delay, and vice versa. Consequently, imposing tighter delay constraints did not force the optimizer to compromise heavily on the cost objective. This results in a Pareto front with very few distinct points, regardless of the ε value.

This objective alignment can often occur in systems where operational efficiency in one echelon inherently supports efficiency in the other (e.g., well-coordinated first-echelon deliveries reducing delays and aiding second-echelon routing). However, this also limits the Pareto front's diversity—important trade-offs may remain hidden.

To justify this claim, we conduct a comparative analysis with randomly generated weights to allow for a broader and less biased exploration of the solution space. More specifically, instead of using the epsilon weights (9), we use randomly generated weights for the objective function of the MIP formulation in (13a) and the allowable delay bounds were relaxed to fall between 10 and 40 minutes. The results are plotted in figure 6 below:

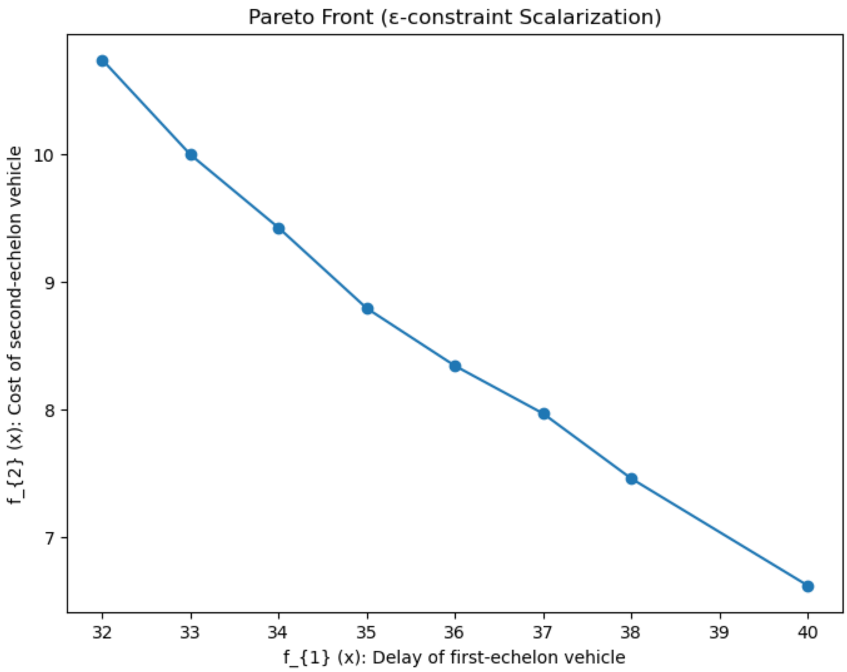

Figure 6: Pareto front generated using ϵ-constraint scalarization with randomized weights of the objective function in the MIP formulation (13a). The plot shows the trade-off between the delay of the first-echelon vehicle and the tour cost of the second-echelon vehicle, revealing a more granular front despite the shorter range of allowable delay bounds.

Here, a far more continuous and well-distributed Pareto front emerges. Each point still represents a non-dominated solution under the ϵ-constraint framework, but the coverage of the trade-off space is much richer. This demonstrates that modifying the weighting logic for generating ε weights (9) of the objective function could significantly improve the diversity and interpretability of the Pareto front.

7 Conclusions and Discussion

In this thesis, we address the critical challenge of optimizing last-mile logistics, with a particular focus on handling multiple, often conflicting objectives. In real-world industrial applications, optimization problems are rarely driven by a single cost criterion. This is especially evident in the distribution of perishable and health-related goods, where both cost efficiency and timely delivery are essential. As a result, there is growing interest in understanding and managing trade-offs between competing objectives in such logistics systems.

To tackle this challenge, we introduce a bi-objective optimization framework for the Combined Second-Echelon Routing Problem (CSERP)—a novel extension of the classical Capacitated Vehicle Routing Problem (CVRP). This model captures the complexity of multimodal, two-echelon last-mile delivery systems. In our formulation, the first echelon is represented by a line-bound public transport vehicle operating on a fixed route and schedule, while the second echelon consists of a single-capacity delivery vehicle responsible for serving customers from designated hand-over points.

The model explicitly addresses two conflicting objectives: minimizing the delay experienced by the first-echelon vehicle, and minimizing the generalized tour cost incurred by the second-echelon delivery vehicle. To capture this trade-off effectively, we developed a cost structure that distinguishes between independent and synchronized cost components. Independent costs reflect the individual travel costs of each vehicle, while synchronized costs capture the time-sensitive coordination required at transfer points.

A mixed-integer programming (MIP) formulation is proposed to solve the CSERP, incorporating constraints related to vehicle capacities, time windows, and routing decisions. The model simulates the interactions between both echelons, ensuring synchronized operations and feasible hand-overs. The delay of the first-echelon vehicle is quantified as the deviation between its scheduled and actual arrival time at the final combined node, while the second objective aggregates both the time-sensitive and travel-related costs associated with the second-echelon tour.

To address the bi-objective nature of the problem, we employ scalarization techniques, specifically the weighted-sum and ϵ-constraint methods, to reduce the problem to a tractable single-objective format. This transformation allows us to systematically explore trade-offs between minimizing delivery delays and reducing routing costs.

Under realistic constraints—such as the assumption of a single, unit-capacity second-echelon vehicle—we show that the problem structure admits a polynomial-time solution using a minimum-weight perfect matching algorithm. This key insight not only enhances computational tractability but also provides a clear interpretation of how hand-over timing and route planning interact. Additionally, the weighted-sum formulation enables flexible sensitivity analysis by reparameterizing objective weights, offering decision-makers a structured way to prioritize specific operational goals.

A detailed numerical example illustrates the proposed methodology and validates the approach. By computing edge weights based on the scalarized objective, we demonstrate how optimal matchings can be efficiently identified using bipartite graph representations.

The ϵ-constraint method proves especially valuable in this context, allowing for direct control over the delay incurred by the first-echelon vehicle while still minimizing overall routing costs. This is particularly useful in time-sensitive logistics scenarios, where some delay is tolerable but must remain within strict operational limits. By fixing upper bounds on delay, the ϵ-constraint method generates a structured Pareto frontier that not only highlights feasible trade-offs but also reveals how sensitive the solution space is to tightening or relaxing delay thresholds.

We examine two extreme cases, unbounded delay ($\varepsilon = \infty$) and zero delay ($\varepsilon = 0$), to showcase the model's flexibility. The former case prioritizes cost minimization with minimal regard for timing, while the latter enforces strict hand-over schedules at the expense of routing flexibility and higher costs. These boundary scenarios help illuminate the operational extremes and offer practical guidance for setting service-level expectations.

We also demonstrated how the CSERP can be reformulated as a mixed-integer matching problem, enabling the use of established optimization software and enhancing solution efficiency.

Between the two scalarization strategies, we found that the ϵ-constraint method was more effective for our intended application. Unlike the weighted-sum method, which can fail to identify non-dominated solutions in certain regions of the Pareto front, the ϵ-constraint method consistently yields such solutions. Moreover, it allows decision-makers to directly influence trade-offs between delay and cost by adjusting ϵ thresholds, making it a more practical and versatile tool. Consequently, the ϵ-constraint method emerged as the preferred approach for solving the CSERP.

Our computational experiments confirmed the practical utility of the proposed model and solution strategies. Using the ϵ-constraint method, we generated a Pareto front that clearly illustrates the trade-off between minimizing delays and reducing tour costs. Despite the model's complexity—including thousands of binary variables and numerous constraints—the solver, aided by heuristics such as the feasibility pump, was able to identify feasible solutions efficiently.

Interestingly, initial experiments with stepwise ϵ values produced only a limited number of non-dominated solutions, indicating a strong correlation between the two objectives: reductions in delay often coincided with reductions in tour cost. This suggests that in well-coordinated logistics systems, improving performance in one echelon can naturally benefit the other. However, this tight alignment also limits the diversity of available trade-offs.

To better explore the solution space, we conducted additional experiments using randomized objective weights, which resulted in a denser and more continuous Pareto front, even within narrow delay bounds. This comparative analysis underscores the importance of using multiple scalarization strategies to fully uncover the range of viable logistics solutions.

In conclusion, this thesis presents a comprehensive framework for tackling bi-objective optimization in two-echelon last-mile logistics. The proposed CSERP model, combined with effective scalarization methods, provides both theoretical insights and practical tools for decision-makers. The results demonstrate not only the model's tractability but also the value of carefully designing objective functions and constraints to expose meaningful trade-offs. Future work could extend this framework to multi-vehicle scenarios, dynamic demand environments, or systems with stochastic travel times, further bridging the gap between academic models and real-world logistics challenges.

References

Affi, M. (2020). General variable neighborhood search approach for solving the electric two-echelon vehicle routing problem. In *2020 International Multi-Conference on: "Organization of Knowledge and Advanced Technologies" (OCTA)*, pages 1–3. IEEE.

Agárdi, A., Kovács, L., and Bányai, T. (2019). Two-echelon vehicle routing problem with recharge stations. *Transport and Telecommunication Journal*, 20:305–317.

Braekers, K., Ramaekers, K., and Van Nieuwenhuyse, I. (2015). The vehicle routing problem: State of the art classification and review. *Computers & Industrial Engineering*, 99:300–313.

Breunig, U., Baldacci, R., Hartl, R. F., and Vidal, T. (2019). The electric two-echelon vehicle routing problem. *Computers & Operations Research*, 103:198–210.

Breunig, U., Hartl, R. F., and Vidal, T. (2017). The two-echelon vehicle routing problem with electric vehicles. In *Simpósio Brasileiro de Pesquisa Operacional*, pages 3797–3807, Blumenau, Brazil.

Caggiani, R., Gatta, G., and Sassi, M. G. (2015a). A survey on two-echelon routing problems. *Computers & Operations Research*, 55:185–199.

Caggiani, R., Gatta, G., and Sassi, M. G. (2015b). A survey on two-echelon routing problems. *Computers & Operations Research*, 55:185–199.

Cattaruzza, D., Absi, N., Feillet, D., and Fernandez, J. G. (2017). Vehicle routing problems for city logistics. *EURO Journal on Transportation and Logistics*, 6(1):51–79.

Clarke, G. and Wright, J. W. (1964). Scheduling of vehicles from a central depot to a number of delivery points. *Operations Research*, 12(4):568–581.

Cook, W. and Rohe, A. (1999). Computing minimum-weight perfect matchings. *INFORMS Journal on Computing*, 11(2):138–148.

Dantzig, G. B. and Ramser, J. H. (1959). The truck dispatching problem. *Management Science*, 6(1):80–91.

Grangeon, J. F., Perboli, G., Tadei, R., and Vigo, D. (2008). The two-echelon capacitated vehicle routing problem. Working Paper.

Guastaroba, G., Speranza, M. G., and Vigo, D. (2016). Intermediate facilities in freight transportation planning: A survey. *Transportation Science*, 50(3):763–789. [Online journal].

Haimes, Y. Y., Lasdon, L. S., and Wismer, D. A. (1971). On a bicriterion formulation of the problems of integrated system identification and system optimization. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-1(3):296–297.

Jie, W., Yang, J., Zhang, M., and Huang, Y. (2019). The two-echelon capacitated electric vehicle routing problem with battery swapping stations: Formulation and efficient methodology. *European Journal of Operational Research*, 272(3):879–904.

Jozefowiez, N., Semet, F., and Talbi, E.-G. (2008a). From single-objective to multi-objective vehicle routing problems: Motivations, case studies, and methods. In Golden, B., Raghavan, S., and Wasil, E., editors, *The Vehicle Routing Problem: Latest Advances and New Challenges*, volume 43 of *Operations Research/Computer Science Interfaces*, pages 445–471. Springer, Boston, MA.

Jozefowiez, N., Semet, F., and Talbi, E.-G. (2008b). Multi-objective vehicle routing problems. *European Journal of Operational Research*, 189:293–309.

Kiba-Janiak, M., Miler, J., Jankowska, A., and Saniuk, A. (2021). Sustainable last mile delivery on e-commerce market in cities from the perspective of various stakeholders. literature review. *Sustainable Cities and Society*, 71. [Online journal].

Marler, R. T. and Arora, J. S. (2004). Survey of multi-objective optimization methods for engineering. *Structural and Multidisciplinary Optimization*, 26(6):369–395.

Mavrotas, G. (2009). Effective implementation of the -constraint method in multi-objective mathematical programming problems. *Applied Mathematics and Computation*, 213(2):455–465.

Moradi, N., Mohammadi Boroujeni, N., Aftabi, N., and Aslani, A. (2024a). Two-echelon electric vehicle routing problem in parcel delivery: A literature review. arXiv:2412.19395 [math.OC].

Moradi, N., Mohammadi Boroujeni, N., Aftabi, N., and Aslani, A. (2024b). Two-echelon electric vehicle routing problem in parcel delivery: A literature review. arXiv:2412.19395 [math.OC].

Park, Y. B. and Koelling, C. P. (1986). A solution of vehicle routing problems in a multiple objective environment. *Engineering Costs and Production Economics*, 10(1):121–132. [Online journal].

Perboli, G., Tadei, R., and Vigo, D. (2011). The two-echelon capacitated vehicle routing problem: Models and math-based heuristics. *Transportation Science*, 45(3):364–380.

Ranieri, L., Digiesi, S., Silvestri, B., and Roccotelli, M. (2018). A review of last mile logistics innovations in an externalities cost reduction vision. *Sustainability*, 10(3):782. [Online journal].

Sandhya and Goel, R. (2018). Multi objective vehicle routing problem: A survey. *Asian Journal of Computer Science and Technology*, 7(3):1–6.

Schiewe, P. et al. (2024). Lintim: An integrated environment for mathematical public transport optimization. Course material.

Schiewe, P. and Stinzendörfer, M. (2024a). The combined second-echelon vehicle routing problem – integrating last-mile deliveries into public transport. Unpublished.

Schiewe, P. and Stinzendörfer, M. (2024b). Optimizing combined tours: The truck-and-cargo-bike case. *OR Spectrum*.

Wang, D. and Zhou, H. (2021). A two-echelon electric vehicle routing problem with time windows and battery swapping stations. *Applied Sciences*, 11(22):10779.

Wang, D., Zhou, H., and Feng, R. (2019). A two-echelon vehicle routing problem involving electric vehicles with time windows. In *Journal of Physics: Conference Series*, volume 1324, page 012071. IOP Publishing.

Wu, Z. and Zhang, J. (2023). A branch-and-price algorithm for two-echelon electric vehicle routing problem. *Complex & Intelligent Systems*, 9:2475–2490.

Zijlstra, T., Galindo Pecin, D., and Spliet, R. (2021). Integrating electric vehicles in the two-echelon vehicle routing problem. *Transportation Science*. Unpublished.