Capacitated vehicle routing problem with a zone tariff

Niklas Tuma, Manuel Ostermeier, Alexander Hübner

Angaben zur Veröffentlichung / Publication details:

Tuma, Niklas, Manuel Ostermeier, and Alexander Hübner. 2026. "Capacitated vehicle routing problem with a zone tariff." *Networks* 87 (2): 210–27.
https://doi.org/10.1002/net.70015.

Networks

WILEY

RESEARCH ARTICLE OPEN ACCESS

Capacitated Vehicle Routing Problem with a Zone Tariff

Niklas Tuma[1] | Manuel Ostermeier[2] | Alexander Hübner[1]

[1]Supply and Value Chain Management, Technical University of Munich, Straubing, Germany | [2]Resilient Operations, University of Augsburg, Augsburg, Germany

Correspondence: Alexander Hübner (alexander.huebner@tum.de)

Received: 10 October 2024 | **Revised:** 7 July 2025 | **Accepted:** 24 September 2025

Funding: The authors received no specific funding for this work.

Keywords: branch-and-check | detour limit | exact approach | logistic services | volume discount

ABSTRACT

High demand and volume fluctuations, constrained transportation capacity, and weak network effects compel retailers to rely on external logistic service providers (LSPs) instead of operating their own transportation fleet. This outsourcing process requires retailers to propose tours to LSPs that then carry out the store deliveries. LSPs typically bill their service using a tariff with a zone structure representing geographical positions. The transportation costs depend on the zones visited and the delivery volume, which is subject to volume discounts. Further, LSPs apply detour limits to prevent excessive travel distances. Zone tariffs are a standard pricing scheme to ease transportation planning in the industry. Despite its high practical relevance, the literature in this field is limited, and a general model and solution approach are lacking. Our work closes this gap by providing the first comprehensive model for the Capacitated Vehicle Routing Problem with a Zone Tariff (C-VRP-ZT) and by developing a generally applicable exact solution method. The developed Branch-and-Check (BAC) framework includes valid inequalities and multiple acceleration techniques. We prove the computational efficiency of our approach using benchmark instances and derive managerial insights. We analyze different tariff structures concerning central characteristics such as the zone layout and order consolidation. Our results show how the tariff characteristics impact the cost of retailers and the margins of LSPs. We further solve a real-world application in retailing and analyze the cost/revenue split between a retailer and an LSP. These results provide insights into the threshold when a retailer should outsource its transportation process.

1 | Introduction

Companies commonly use logistic service providers (LSPs) to master their distribution processes, transferring transportation tasks to external partners. The market size for these outsourced transportation services is substantial: LSPs are expected to generate worldwide revenue of 1.44 trillion $ in 2024 [1]. They provide services for 90% of the Fortune 500 companies in the United States [2], including leading market players such as Apple, Nvidia, and Tesla. Also, most retailers depend on LSPs to master their distribution due to the arising logistic complexity of their operations. Employing LSPs for their operations offers retailers many advantages compared with maintaining an in-house fleet. Retailers frequently encounter substantial demand fluctuations, particularly during seasonal peaks like Christmas, which leads to significant variations in transportation volumes. Weekly promotions and changing assortments with diverse transportation needs and volumes (e.g., garden supplies vs. packaged electronic equipment) add to the complexity of the distribution and necessitate highly flexible operations. These circumstances induce many retailers to outsource their entire transportation operations. Examples are leading retailers such as Aldi, Lidl, or MediaMarkt-Saturn. However, even retailers operating their own fleet, such as Walmart or Tesco, use LSPs to cover peaks during

https://doi.org/10.1002/net.70015

high-demand seasons. The benefits of resorting to LSPs are obvious: LSPs offer large and flexible networks, pooling fluctuating transportation needs across their customers.

Retailers usually supply their network from a distribution center (DC) via less than truckload (LTL) deliveries. Outsourcing the corresponding transportation tasks to an LSP constitutes a central planning problem of a retailer. Addressing the arising challenges requires seamless coordination between a retailer and the cooperating LSP. The operational planning generally follows two steps. First, the retailer consolidates deliveries for multiple stores to achieve economies of scale in transportation. The consolidated store orders are then suggested as combined deliveries to one LSP. Second, the LSP assesses the feasibility and cost-effectiveness of the proposed tours and handles the actual transportation once approved. The LSP typically charges the transportation services using *zone-based tariffs* to ease the planning and coordination with the retailer. Tariff zones are the standard within municipal public transport, where the number of traversed zones during a journey defines the prices [see e.g., 3]. Zone-based tariffs are further common in parcel delivery services [see e.g., 4, 5]. Similarly, LSPs utilize geographical zones to represent travel distances in goods transportation [see e.g., 6, 7], and charge based on the zone(s) visited on a tour. Zone-based tariffs considerably simplify rate calculations by offering a predefined price structure instead of pricing actual travel times and distances. This structure enables an eased prediction of transportation costs, streamlines distribution planning and coordination for LSPs and retailers, and ultimately reduces complexity for both parties. In addition to the zones, the second dimension of the tariff structure is volume discounts for increasing delivery quantities. Tours consisting of multiple stores are then charged for the farthest zone visited on the tour and the total transportation load. Although zone tariffs simplify the process by eliminating the need to solve a complex routing problem, it remains crucial for the retailer to establish meaningful store clusters for each tour. When minimizing transportation costs using the tariff structure, the retailer can only combine stores on a tour that adheres to capacity limits and specific travel distance constraints between stores. An LSP may otherwise reject the retailer's proposals.

Despite the widespread practical significance of zone tariffs across industries, only a few related applications are discussed in the literature. The existing works primarily address very case-specific distribution problems involving zone-based tariffs. Most strikingly, there is no comprehensive Mixed-Integer Program (MIP) of the transportation problem with zone-based tariffs of external providers, where a retailer must propose feasible tours to be executed by an LSP. The pricing schemes and practical constraints derived from the industry applications have not yet materialized in efficient decision support systems. Moreover, the industry strives to support decisions with advanced models and solution approaches. The prevailing planning often relies on the planners' experience or simple rule-of-thumb methods, while advanced modeling and optimization are lacking [see, e.g., 8].

Our work contributes to bridging this literature gap by providing a comprehensive MIP that formalizes the retailer's business problem of outsourcing transportation tasks to an LSP. As such, it offers a practical solution to a well-known industry challenge and addresses a real-world need. Drawing on the tariff structures and planning methodologies observed in practice, we formalize the Capacitated Vehicle Routing Problem with a Zone Tariff (C-VRP-ZT) and propose an efficient exact approach. We develop a Branch-and-Check (BAC) framework that decomposes the C-VRP-ZT into a master problem for assigning stores to tours and a subproblem for obtaining feasible store combinations respecting tour constraints. The BAC enables us to solve instances of practically relevant size and to exemplify the planning within an application in the retail industry. Our numerical studies provide insights on how zone-based tariffs impact transportation planning compared to the classical distance-based Vehicle Routing Problems (VRPs); how retailers can make use of the tariff structures to minimize cost; and how different tariff features affect the cost of retailers and margins of LSPs.

We organize the remainder of our work as follows: Section 2 details the planning problem and discusses related literature. Section 3 introduces the C-VRP-ZT and develops a problem-specific BAC approach. We then analyze the numerical performance of the novel approach, provide managerial insights on tariff structures, and solve a real-world application (Section 4). Section 5 summarizes our findings.

2 | Problem Description and Related Literature

Zone tariffs for goods transportation have not been comprehensively introduced in the prevailing literature. To reflect on the planning problem faced in the industry, we will first delineate the general problem setting, discuss the key aspects of the zone tariff scheme, and then outline the coordination process and the implied planning discrepancy. Section 2.1 therefore develops the framework for analyzing related literature and derives the research gap in Section 2.2. Our work focuses on the application in retail, but the setting with LSPs for transportation services generally applies across industries.

2.1 | Problem Setting and Planning Process

The majority of retailers, in particular in non-food applications, rely on external carriers and supply most of their stores via DCs [9]. Deliveries are usually executed on joint store tours, regardless of whether an own or external fleet of trucks is used. The trucks' capacities limit the store order sizes, while stores mostly order LTL. When planning picking capacities in DCs and replenishment processes in stores, retailers establish centralized guidelines for order and delivery schedules at a mid-term tactical level [10, 11]. This means stores can order on weekdays according to the assigned delivery schedules of the retailer. Orders must be placed at least 1-2 days in advance to align DC operations and to define delivery tours based on the actual order sizes and delivery volumes. To achieve economies of scale, retailers often partner with a specific LSP for each DC. The daily operational transportation planning is carried out by logistics planners at the DC. They determine which stores are supplied together on a tour without defining the actual sequence of the tour and transfer the tour requests for execution to the LSP. At this point, the LSP only provides the tariff structure and available capacities.

TABLE 1 | Example of a tariff scheme: Costs per zone and total load.

		Farthest visited zone of tour				
		Zone 1	**Zone 2**	**Zone 3**	**Zone 4**	**…**
Total load on tour	…	…	…	…	…	
	5 Units	€25.00	€27.50	€31.00	€34.50	…
	6 Units	€23.33	€26.50	€29.83	€33.33	…
	7 Units	€22.14	€25.79	€29.00	€32.50	…
	8 Units	€21.25	€25.25	€28.38	€31.88	…
	9 Units	€20.56	€24.83	€27.89	€31.39	…
	10 Units	€20.00	€24.50	€27.50	€31.00	…
	…	…	…	…	…	…

Tariff structure. The LSPs offer predetermined tariffs to charge for their services, such that retailers can instantly determine the cost of their transportation requests. Table 1 illustrates a typical tariff scheme where the tour costs depend on two components: the *zone* of the farthest store visited and the *total load*. An exemplary tour with a store from Zone 1 and a second store from Zone 2, both ordering 5 units, results in 24.50 € per load unit (see Table 1; Zone 2 as farthest zone and a total of 10 load units) and a total cost of 245 €(= 10 · 24.50€). Comparing the total cost of two separate shipments for the same stores amounts to 262.50 €. However, not all store combinations are economically meaningful. For example, adding stores of lower zones to tours with a higher zone increases the total load and entitles them to higher volume discounts. Still, the combination may lead to higher total costs than separate shipments. Assuming a tour with one store from Zone 1 and a second store from Zone 4, both ordering 5 units, results in a total cost of 310 €, while separate shipments amount to 297.50 €. These examples show the retailers' cost trade-offs when defining the delivery tours for the LSP. Retailers must find meaningful store combinations that leverage the volume discounts and reduce distance costs due to the different tariff zones, resulting in minimal total transportation costs.

Our tour examples show that the total cost depends on the *(1) zone layout* and the *(2) volume discounts*. Moreover, to ensure feasible tours, there are *(3) LSP restrictions* for travel distance and vehicle capacity. We will further detail these aspects of the tariff structure.

(1) Zone layout. The zone layout defines the number of zones and their geographical size and form. Each zone is a geographical area containing all stores with equal transportation costs. The typical layout is concentric, and the price of a zone is determined by its distance from a fixed central point, typically the source (e.g., the DC) [12]. Other geographical features like ZIP codes or city districts may be used to further determine the zone layout [see, e.g., 13]. All zoning approaches have in common that each customer (in our case, store) is uniquely assigned to one zone.

(2) Volume discounts. To incentivize retailers to combine multiple shipments within their requests, LSPs offer volume discounts, sharing the cost savings from higher total loads on a tour. Table 1 shows this effect as the costs per load decrease with an increasing total load. The price for each zone-load combination is based on (a) a fixed base cost for each zone independent of the load and (b) a variable cost. The decreasing variable costs award volume discounts for higher load sizes. However, volume discounts result in a non-linear cost structure and complicate the decision-making. One way to address the non-linearity is the definition of piece-wise and step-wise linear functions or discretized tariffs when designing the pricing scheme [see e.g., 8, 14].

(3) LSP restrictions. The tour definition is first restricted by the trucks' capacity limits. Considering the economies of scale of volume discounts, retailers are incentivized to increase capacity utilization. Alongside the capacity restrictions, LSPs apply so-called *detour limits* as zone-based tariffs potentially allow extensive tours: assigning stores from the same zone but with significantly different geographical locations to a single tour does not increase the retailer's costs under the zone tariff but may lead to substantially longer driving distances for the LSP. Figure 1 illustrates a stylized example with two stores in Zone 2 but on different sides of the DC. Combining both stores on a tour results in a low cost for the retailer at the expense of disproportionate driving distances for the LSP. Consequently, the LSPs restrict the in-tour driving distances by limiting the allowed detour. The detour is also referred to as out-of-route distance or ratio [see, e.g., 15]. The retailer does not need to consider the return of the truck to its DC for determining the detour. Only the driving distances from the DC to the stores are relevant since trucks may return to the LSP's depot after finishing the tour or complete transportation services for other customers.

Figure 1 shows a detour calculation when combining a store from Zone 1 with a store from Zone 2. The absolute detour is defined as the difference between the shortest travel distance of a tour (in the example 146 km + 368 km = 514 km) and the direct distance from the DC to the farthest store on the tour (477 km), that is, the absolute detour amounts to 514 km − 477 km = 37 km. The detour limit defines the corresponding upper bound of the allowed detour on a delivery tour. It is a constant and equal for all tours (e.g., 100 km). The absolute detour limit is standard in real-world applications [see also 12, 15]. The relative detour limit is an alternative. It is defined as the shortest travel distance of a tour divided by the direct distance from DC to the farthest store on the tour, where the respective limit is a constant factor (e.g., 1.2). Both variants require knowing the travel distance when proposing joint tours so that the detour limit is maintained. This requires solving a Traveling Salesperson Problem (TSP) to sequence the stores on the tour.

Coordination and planning discrepancy. The outsourcing of the distribution process to an LSP combined with the zone tariff requires the concerted interaction of the retailer and the LSP. The core responsibility for the planning rests with the retailer's

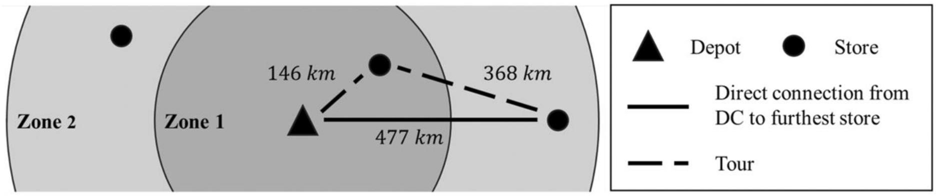

FIGURE 1 | Calculation of absolute detour (example).

planners. They are responsible for suggesting economically meaningful tours to the LSP. The major challenge is to determine suitable store tours (i.e., cluster store orders to obtain low transportation costs) without actually defining the specific travel sequences, as the LSP executes the tours according to its own planning, and detailed routing from the retailer is not required. This constitutes a discrepancy in the retailer's planning: While the retailer aims to maximize cost savings by leveraging the zone structure and volume discounts, the tours must adhere to vehicle capacity constraints and the detour limits—and the latter requires detailed distance information of the tours planned. Not adhering to the constraints specified by the LSP leads to the rejection of tours and increased planning efforts and costs for the retailer. Defining economical and feasible tour requests is decisive for retailers to minimize their distribution costs. While the tariff clearly defines the tour costs for retailers, the actual costs for the LSP are unknown for the retailer, as the LSP may consolidate shipments of multiple customers (e.g., different retailers) to increase capacity utilization. This drives the LSP's margins and highlights the benefits of offering the zone tariffs to customers.

2.2 | Related Literature

This section identifies and analyzes the relevant contributions regarding zone tariffs in goods transportation, focusing on essential planning characteristics for real-world applications (e.g., detour limits and volume discounts). Despite their relevance in the industry, where zone tariffs are a standard pricing scheme to facilitate transportation planning, research on zone-based tour planning is scarce, highlighting a valuable opportunity for further academic exploration. Common across all identified papers is a setting where a company determines capacitated tours that are proposed to and carried out by LSPs.

Ceselli et al. [13] are the first to address a zone-based transportation problem, where customers belong to several hierarchical zones that are defined by geographical features (e.g., nation, region, district, etc.). In their application, zone costs are determined by multiple characteristics (e.g., load, number of stops, value, weight, etc.), and the highest cost across all stops determines the actual tour cost. Volume discounts are represented by a piece-wise linear function. The transportation problem further includes a heterogeneous fleet, multiple depots and carriers, time windows, and an outsourcing option. The authors propose a Branch-and-Price (BP) approach involving a Set Partitioning Model (SPM) to solve a case study. Tours are generated within the subproblem of their BP framework. Ceschia et al. [16] present a load- and zone-dependent tariff function next to other cost functions to evaluate a transport problem with time windows, carrier-dependent costs, and the option to use a second delivery mode for unscheduled orders. Further extensions include a heterogeneous fleet and a multi-period horizon, while no formal model of the decision problem is presented. They apply a Tabu Search (TS) framework to solve benchmark instances and a case study. Tordecilla et al. [17] cluster the delivery area into zones and apply a tariff that prices tours dependent on the most expensive zone, combined with the number of stops. They do not consider a volume discount and develop a specialized heuristic (SH) to solve a case in the agri-food industry, constituting a rich transportation problem with multi-compartment vehicles, a heterogeneous fleet, and stop priorities. The authors do not provide a comprehensive mathematical representation of their problem variant.

The works above do not apply any kind of detour restrictions, which substantially reduces the complexity. The first contribution considering these restrictions in our context is Soleilhac et al. [12]. Their zone-based tariff is based on concentric zones around the DC and an absolute detour limit for deliveries. The cost structure depends on the farthest crossed zone and additionally includes stop costs, while discounts are neglected for truck tours. The retail transportation problem comprises multiple carriers, a heterogeneous fleet, and a second mode for single shipments. The authors develop a Large Neighborhood Search (LNS) to solve benchmark and case instances and provide an SPM for the tour selection and outsourcing decisions. Their approach does not rely on a full mathematical representation of the problem. Finally, Tuma et al. [8] consider a tariff variant with concentric zones around the DC, volume discounts for delivery tours, and a relative detour limit. The arising transportation problem is defined by two delivery modes, a heterogeneous fleet, open routes, and multiple carriers. They propose a software tool based on a Set Partitioning Algorithm (SPA) that is designed for their industry application. Due to the constrained problem structure of the specific application, selecting optimal tours is possible for a restricted set of pre-computed tours using a tree-based algorithm. Although the problem is well aligned with the practical business context, a formal MIP of the entire decision problem is not provided.

Literature gap. Table 2 summarizes the related publications dealing with zone-based tariffs in goods transportation. The fact that all related contributions are driven by real-world problems underlines the proximity to industrial applications but also indicates the lack of theoretical advancements so far. The existing publications share a common planning context, where a shipper proposes delivery plans with order consolidation and limited

TABLE 2 | Overview of related literature with zone-based tariffs.

Publication	Tariff structure: Criteria[c]	(1) Zone layout	(2) Vol. discounts	(3) Detour limit	Decision model[a]	Method	Instance type[b]
Ceselli et al. [13]	Multiple	Geograph.	✓	—	(✓)	BP	Case
Ceschia et al. [16]	Load	Not spec.	✓	—	—	TS	Case & benchmark
Tordecilla et al. [17]	Stop	Not spec.	—	—	—	SH	Case
Soleilhac et al. [12]	Stop	Concentric	—	Absolute	—	LNS	Case & benchmark
Tuma et al. [8]	Load	Concentric	✓	Relative	—	SPA	Case
This paper	Load	Concentric	✓	Absolute	✓	BAC	Case & benchmark

[a]Formal representation/mathematical formulation of the entire decision problem (e.g., as MIP), including tour generation and tour selection.
[b]Instances solved, that is, (application) case: real-world case; benchmark: reproducible and generally available benchmark instances.
[c]Components additional to zone-based criteria; Multiple includes distance, number of pallets, weight, volume, value & stops.

truck capacity to an LSP using a tariff structure. Furthermore, they apply a zone-based structure with different zone layouts consisting of zones and additional pricing factors (e.g., stops, loads) to determine the total transportation cost. The three contributions considering total load in their objective also apply volume discounts as part of the tariff. The two studies that analyze concentric zones further apply detour limits.

Remarkably, none of the existing contributions present a compact model formulation that captures the full scope of the decision problem and provides a full mathematical representation of the business problem. Such a formulation would encompass the entire VRP in one comprehensive model, including both the assignment of customers to tours and the sequencing of customers within each tour. As existing works solve case-specific applications, they focus on developing targeted solution approaches. There is a mixture of heuristics and exact approaches. The exact approaches are applied to application-specific, rich problems—solving highly constrained variants with particular cost structures of the respective industry problem. This limits the comparability of these solution approaches and does not prove their applicability to case-independent problem settings.

There is no contribution addressing all three industry-relevant tariff features ((1) zone layout, (2) volume discounts, and (3) LSP constraints)), providing a MIP of the complete decision problem and solving reproducible benchmark instances. This also limits the managerial implications of using zone-based tariffs for practitioners and researchers alike. We close this gap and develop the first MIP for the entire C-VRP-ZT with a zone tariff, detour limit, and volume discounts. We further solve benchmark and case-specific instances of the novel problem with an efficient exact approach.

3 | Mathematical Model and Solution Approach

3.1 | Formal Definition

This section develops the formal representation of the C-VRP-ZT verbalized in Section 2. The retailer aims to minimize transportation costs within the scope of the zone-based tariff by defining delivery tours that are executed by an LSP while adhering to detour and capacity limitations. The C-VRP-ZT ensures the cost-optimal tour generation with volume discounts for retailers and the tour acceptance by LSPs. Table 3 introduces the notation.

TABLE 3 | Notation.

Index sets	
I	Set of stores excluding the DC (node 0)
L	Set of possible vehicle loads
T	Set of vehicle tours
Z	Set of tariff zones
Parameters	
$c_{z,l}$	Tariff costs for zone z with total load l
$d_{i,j}$	Distance from location i to location j, $i, j \in I \cup \{0\}$
γ_i	Demand of store i
r_i	Tariff zone of store i
Δ	Absolute detour limit
Q	Vehicle capacity
Variables	
u_i	Continuous: tour distance to reach store i
v_t	Continuous: total distance of open tour t
$n_{i,j}$	Binary: 1, if store i is visited on a tour before store j; 0, otherwise
$w_{i,t}$	Binary: 1, if store i is selected for the detour limit check; 0, otherwise
$x_{i,t}$	Binary: 1, if store i is included in tour t; 0, otherwise
$y_{t,z,l}$	Binary: 1, if the tariff for tour t is defined by zone z and load l; 0, otherwise

The C-VRP-ZT is defined on a directed graph $G(N, E)$, with a set of nodes N and edges E. The set $N = I \cup \{0\}$ includes the DC (node 0) and the set of stores I, $i \in I$. The set $E = \{(i, j) : i, j \in N, i \neq j\}$ consists of all edges between the nodes N. The edges are associated with the travel distances $d_{i,j}$, fulfilling the triangle inequality. All tours originate at the retailer's DC, but due to the zone tariff and the use of an LSP, there is no return to this DC. Each tour, therefore, constitutes an Open TSP (O-TSP) starting at node 0 and ending at any store $i, i \in I$. The LSP operates a set of vehicles that is sufficiently large to supply all stores (i.e., the number of vehicles is not limited). The LSP uses homogeneous vehicles with capacity Q to supply the stores via delivery tours $t \in T$, where each tour carries a discrete number of load units $l \in L$, with $L = \{1, \dots, Q\}$. Stores have a non-negative demand γ_i, and each store is allocated to one specific tariff zone, denoted by r_i. The set Z with $z \in Z$ summarizes all tariff zones. We denote the absolute detour limit with Δ. The tour costs $c_{z,l}$ depends on the farthest zone z visited and the total load l carried on the tour.

We introduce the following variables: The continuous variable u_i represents the travel distance of a tour to reach store i, and v_t denotes the travel distance of the entire open tour t. We further define four binary variables: $x_{i,t}$ resembles the assignment decision of stores to tours, and $w_{i,t}$ sets the distance limit for a tour by indicating the farthest away store i on tour t. The variable $n_{i,j}$ defines the precedence relation of all stores on a tour. Finally, $y_{t,z,l}$ indicates the tariff zone z and load l that applies for tour t. We formulate the C-VRP-ZT with Equations (1) to (17).

$$\text{Minimize} \quad \sum_{t\in T}\sum_{z\in Z}\sum_{l\in L} c_{z,l} \cdot y_{t,z,l} \tag{1}$$

subject to

$$\sum_{t\in T} x_{i,t} = 1 \qquad \forall i \in I \tag{2}$$

$$\sum_{z\in Z}\sum_{l\in L} y_{t,z,l} \le 1 \qquad \forall t \in T \tag{3}$$

$$\sum_{i\in I} \gamma_i \cdot x_{i,t} \le \sum_{z\in Z}\sum_{l\in L} l \cdot y_{t,z,l} \qquad \forall t \in T \tag{4}$$

$$r_i \cdot x_{i,t} \le \sum_{z\in Z}\sum_{l\in L} z \cdot y_{t,z,l} \qquad \forall t \in T,\ i \in I \tag{5}$$

$$n_{i,j} + n_{j,i} \ge x_{i,t} + x_{j,t} - 1 \qquad \forall t \in T,\ i, j \in I,\ i < j \tag{6}$$

$$u_i \ge d_{0,i} \qquad \forall i \in I \tag{7}$$

$$u_j \ge u_i + d_{i,j} - M(1 - n_{i,j}) \qquad \forall i, j \in I,\ i \ne j \tag{8}$$

$$v_t \ge u_i - M(1 - x_{i,t}) \qquad \forall t \in T,\ i \in I \tag{9}$$

$$v_t - \sum_{i\in I} d_{0,i} \cdot w_{i,t} \le \Delta \qquad \forall t \in T \tag{10}$$

$$w_{i,t} \le x_{i,t} \qquad \forall t \in T,\ i \in I \tag{11}$$

$$\sum_{i\in I} w_{i,t} \le 1 \qquad \forall t \in T \tag{12}$$

$$u_i \in \mathbb{R}_0^+ \qquad \forall i \in I \tag{13}$$

$$v_t \in \mathbb{R}_0^+ \qquad \forall t \in T \tag{14}$$

$$n_{i,j} \in \{0,1\} \qquad \forall i, j \in I,\ i \ne j \tag{15}$$

$$w_{i,t},\ x_{i,t} \in \{0,1\} \qquad \forall i \in I,\ t \in T \tag{16}$$

$$y_{t,z,l} \in \{0,1\} \qquad \forall t \in T,\ z \in Z,\ l \in L. \tag{17}$$

The objective function (1) minimizes the total cost of all operated delivery tours t, considering the most distant zone z visited and the total load l carried on each tour. Equations (2) ensure that each store i is allocated to exactly one tour t. Constraints (3) specify that at most one active tariff can be selected for each tour, while Inequalities (4) and (5) ensure that the selected tariff must adhere to the total load l and the highest zone z. Constraints (6) define the precedence relations of two stores i and j assigned to the same tour, that is, if the tour visits i before or after j. Constraints (7) specify that the current tour distance at store i must be greater than or equal to the respective direct distance from node 0. Constraints (8) then determine the current tour distance when arriving at store j, where M is a sufficiently large number. Similarly, Constraints (9) determine the total tour distance of each tour. Constraints (10–12) ensure the detour limit adherence. The first ensures that the detour does not exceed the limit Δ. Constraints (11) and (12) define the corresponding auxiliary variable $w_{i,t}$. Together, the constraints determine a store with sufficient distance from the depot such that the detour limit is adhered to. Please note that for checking the tour feasibility, it is not necessarily required to set $w_{i,t} = 1$ for the farthest store from the depot. In other words, if Constraints (10) are fulfilled by selecting a nearer store, they will also be fulfilled by selecting the farthest store. The variable can be set to one at most for one store i on tour t and only if the store is on this tour. Constraints (13–17) define the variable domains.

3.2 | Valid Inequalities

The C-VRP-ZT is NP-hard since the routing constraints resemble a generalization of the TSP. We therefore propose valid inequalities to strengthen the problem formulation. The problem is prone to symmetric solutions since we consider homogeneous vehicles and a feasible tour can be carried out by any vehicle. We introduce the symmetry-breaking Constraints (18) proposed by [18] to limit the number of parallel solutions. These valid inequalities ensure that the first store assigned to tour t has a lower or equal index than the first store assigned to tour $t+1$. This logic also results in the use of tours in ascending order. Further, Constraints (19) ensure that no tariff can be set if no stores are assigned to tour t, and (20) limits the total load carried on a tour to the vehicle capacity Q. Finally, we define the parameter o as the lower bound of active tours required. We determine o by solving a bin packing problem to minimize the number of tours while adhering to the capacity limit and the store assignment constraints. The parameter o allows us to formulate Constraints (21) ensuring a minimum number of active tours $y_{t,z,l} = 1$. Together with Constraints (2–17), the valid inequalities result in an extended MIP formulation for the C-VRP-ZT that represents the identical decision problem strengthened by the introduced inequalities (18–21).

$$x_{j,t+1} \le \sum_{i=1}^{m} x_{i,t} \qquad \forall j, m \in I,\ j \le m,\ t \in \{1, \dots, |T| - 1\} \tag{18}$$

$$\sum_{z\in Z}\sum_{l\in L} y_{t,z,l} \le \sum_{i\in I} x_{i,t} \qquad \forall t \in T \tag{19}$$

$$\sum_{i\in I} \gamma_i \cdot x_{i,t} \le Q \qquad \forall t \in T \tag{20}$$

$$\sum_{t\in T}\sum_{z\in Z}\sum_{l\in L} y_{t,z,l} \ge o. \tag{21}$$

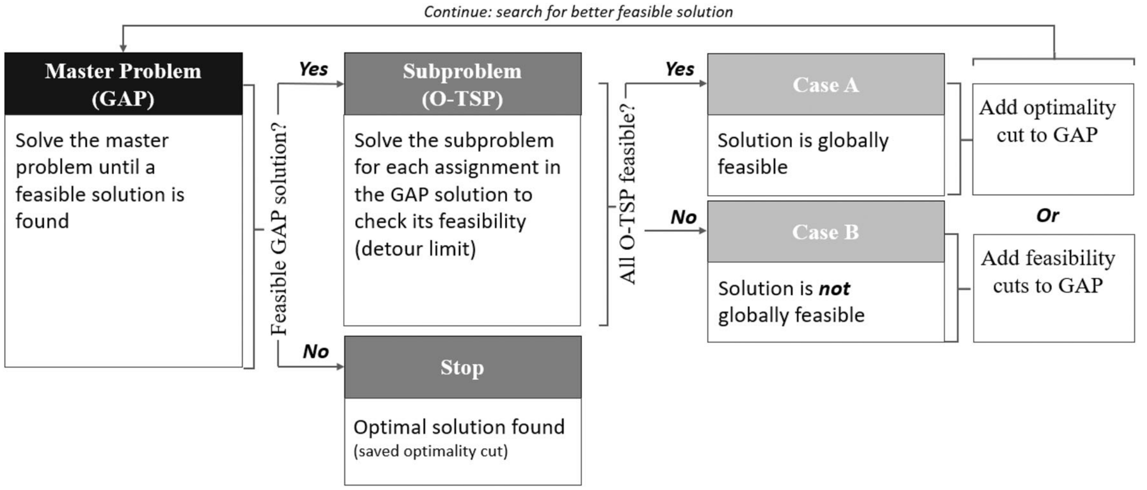

FIGURE 2 | Illustration of one iteration of the BAC algorithm.

3.3 | Branch-and-Check Solution Algorithm

The C-VRP-ZT assigns stores to tours, while the tours built must adhere to detour and capacity constraints. It therefore combines two problem types: the Objective Function (1) and Constraints (2–5) form a Generalized Assignment Problem (GAP), and Equations (6–12) represent a generalization of an O-TSP for multiple tours. We exploit this structure and separate the clustering and routing [see, e.g., 19] by decomposing the problem into a master problem (GAP) and subproblem (O-TSP). In our context, the GAP master problem constitutes the challenging part comprising the complete set of stores. In contrast, the O-TSP subproblem is computationally less demanding due to the limited tour sizes enforced by Δ and Q, thus considering only a small subset of stores per tour. We therefore develop a BAC approach to solve the decomposed problem. The BAC is well-suited for similar problems [see, e.g., 20, 21]. It solves the subproblem for every feasible solution found by the master problem. This implies that the GAP does not need to be solved to optimality in each iteration but only until a new feasible solution is found (see Figure 2). The BAC was first proposed by [18] and generalizes the well-known Benders decomposition [22] by solving subproblems of any form, including MIPs. We further refer to [23] for a detailed review of Benders decomposition approaches.

Figure 2 summarizes the BAC framework developed, illustrating one iteration. We first solve the master problem (GAP) until we find a feasible assignment of stores to tours while respecting the capacity constraints. However, this assignment and the resulting tours may violate the detour limit. We consequently check the feasibility by solving the subproblem for each tour assignment determined by the master problem. If all tours adhere to the detour limit (Case A in Figure 2), the assignment constitutes a globally feasible solution for the C-VRP-ZT. The globally feasible solution is then saved, and a respective optimality cut is added to the master problem. The search for a lower-cost assignment within the master problem continues. Otherwise (Case B), we forbid infeasible tours found within the subproblem by adding feasibility cuts to the master problem. Once the master problem cannot find any feasible assignment, the best-saved solution (optimality cut) is the optimal solution of the C-VRP-ZT. To accelerate the search, we additionally propose problem-specific acceleration techniques (e.g., an upper bound of tours; see Section 3.3.4) used within the BAC.

3.3.1 | GAP Master Problem

The GAP ultimately determines a cost-optimal assignment of stores to tours. This assignment adheres to the vehicle capacity and customer assignment constraints but excludes the routing constraints and, in particular, the detour limit. The GAP is formulated as follows.

$$\text{Minimize} \quad \sum_{t\in T}\sum_{z\in Z}\sum_{l\in L} c_{z,l}\cdot y_{t,z,l} \tag{1}$$

subject to

$$\sum_{t\in T} x_{i,t} = 1 \qquad \forall i \in I \tag{2}$$

$$\sum_{z\in Z}\sum_{l\in L} y_{t,z,l} \le 1 \qquad \forall t \in T \tag{3}$$

$$\sum_{i\in I} \gamma_i \cdot x_{i,t} \le \sum_{z\in Z}\sum_{l\in L} l \cdot y_{t,z,l} \qquad \forall t \in T \tag{4}$$

$$r_i \cdot x_{i,t} \le \sum_{z\in Z}\sum_{l\in L} z \cdot y_{t,z,l} \qquad \forall t \in T,\ i \in I \tag{5}$$

$$y_{t,z,l} \in \{0,1\} \qquad \forall t \in T,\ z \in Z,\ l \in L \tag{17}$$

$$x_{j,t+1} \le \sum_{i=1}^{m} x_{i,t} \qquad \forall j,\ m \in I,\ j \le m,\ t \in \{1,\dots,|T|-1\} \tag{18}$$

$$\sum_{z\in Z}\sum_{l\in L} y_{t,z,l} \le \sum_{i\in I} x_{i,t} \qquad \forall t \in T \tag{19}$$

$$\sum_{i\in I} \gamma_i \cdot x_{i,t} \le Q \qquad \forall t \in T \tag{20}$$

$$\sum_{t\in T}\sum_{z\in Z}\sum_{l\in L} y_{t,z,l} \ge o \qquad \forall t \in T \tag{21}$$

$$\text{Feasibility/optimality cuts} \qquad \text{from subproblem} \tag{22}$$

$$x_{i,t} \in \{0,1\} \qquad \forall i \in I,\ t \in T. \tag{23}$$

This formulation is identical to the C-VRP-ZT except for the routing-related Constraints (6–12). We further add the symmetry-breaking Constraints (18), valid inequalities (19–21), and the feasibility/optimality cuts from the subproblem in (22) (see Section 3.3.2 below). The GAP master problem within the BAC is solved to obtain a feasible solution that clusters all stores to tours. These assignments are then input for the subproblem to determine the routes of tours.

3.3.2 | O-TSP Subproblem

The subproblem checks the global feasibility of every GAP solution found (tour assignment). The global feasibility requires adhering to the absolute detour. This necessitates solving the shortest path problem in the form of an O-TSP, that is, the routing part of the C-VRP-ZT that has been excluded in the master problem. As we solve the O-TSP for each tour determined, we introduce the set T^{GAP} representing the store assignments of the current GAP solution, and $I_t \subseteq I$ indicating the set of stores i assigned to tour $t \in T^{\mathrm{GAP}}$. We adapt the routing constraints of the C-VRP-ZT to a single tour such that the variables v and w_i do not include the tour index t. The O-TSP then minimizes the total distance of a tour t, indicated by d_t^*:

$$\text{Minimize} \qquad v = d_t^* \tag{24}$$

subject to

$$n_{i,j} + n_{j,i} \geq 1 \qquad \forall i,j \in I_t,\ i<j \tag{25}$$

$$u_i \geq d_{0,i} \qquad \forall i \in I_t \tag{26}$$

$$u_j \geq u_i + d_{i,j} - M(1-n_{i,j}) \qquad \forall i,j \in I_t,\ i \neq j \tag{27}$$

$$v \geq u_i \qquad \forall i \in I_t \tag{28}$$

$$v - \sum_{i \in I_t} d_{0i} \cdot w_i \leq \Delta \tag{29}$$

$$\sum_{i \in I_t} w_i \leq 1 \tag{30}$$

$$v \in \mathbb{R}_0^+ \tag{31}$$

$$u_i \in \mathbb{R}_0^+ \qquad \forall i \in I_t \tag{32}$$

$$n_{i,j} \in \{0,1\} \qquad \forall i,j \in I_t,\ i \neq j \tag{33}$$

$$w_i \in \{0,1\} \qquad \forall i \in I_t. \tag{34}$$

Constraints (25–34) define the routing constraints for a single tour analogously to the C-VRP-ZT. As the store assignments of the GAP are input to the subproblem and represented by the sets I_t, the assignment variables $x_{i,t}$ are not required. The remaining formulation of the subproblem is conceptually identical to the C-VRP-ZT. In our BAC framework, we apply a dynamic programming algorithm to solve the O-TSP subproblems efficiently. Specifically, we adapt the algorithm of [24] to open tours.

Solving the subproblem for each individual GAP tour assignment results in two possible cases:

A. The subproblem is feasible for all tours (Case A, Figure 1), and the objective value corresponding to the obtained tour assignments (GAP solution) is set as an upper bound by adding an optimality cut to the master problem.

B. The subproblem is *not* feasible for at least one tour t, $t \in T^{\mathrm{GAP}}$ (Case B, Figure 1): The corresponding store assignments I_t are added to the set of infeasible tours (store assignments) S. The set S comprises all store assignments found by the master problem that violate the detour limit and are globally infeasible for the C-VRP-ZT. The generation of the tours in S is forbidden in the master problem using feasibility cuts (see below).

3.3.3 | Feasibility Cuts

This section details Case (B) when tour assignments violate the detour limit and feasibility cuts are required. The feasibility cuts must prevent the generation of globally infeasible tours, while they must not prevent the generation of any globally feasible tour. We propose two types of feasibility cuts. The first type of cuts, referred to as *(B1) detour cuts*, are problem-specific cuts that extend the classical combinatorial "no-good cuts" initially introduced by [25]. The second type, denoted as *(B2) distance cuts*, further enhances our approach.

(B1) Detour cuts. The intention of the cuts is to forbid an infeasible tour and, if possible, further infeasible supersets of the tour. A tour superset t' here defines an extended tour that comprises the set of stores of the tour t and additionally includes further stores, with $I_t \subset I_{t'}$ and $|I_{t'}| > |I_t|$. To derive the problem-specific cuts, we first discuss the special properties of the cuts by providing two propositions and then formulate the actual detour cuts.

Proposition 1. *Let t be any tour provided by the master problem ($t \in T^{\mathrm{GAP}}$). The absolute detour δ_t of tour t with cardinality $|I_t|$ can be higher than the absolute detour of a tour superset t' of tour t, with cardinality $|I_{t'}| > |I_t|$.*

Proof. We prove Proposition 1 by contradiction. Let us assume that the detour always increases when adding stores to a tour. Let us further consider the DC (node 0) at coordinates (0,0) and two tours t_1 and t_2 represented by store coordinates: $I_{t_1} = \{(8,1),(8,-1)\}$, and $I_{t_2} = \{(8,1),(8,-1),(9,0)\}$. The extended tour t_2 is then a superset of t_1. We assume Euclidean distances between the locations and calculate the detour of both tours (i.e., solving the O-TSP subproblem) by $\delta_t = d_t^* - max_{i \in I_t}(d_{0,i})$, with d_t^* representing the minimum travel distance of tour t (see Equation (24)). The DC is the starting point of the tour. The given example results in the following distances and detours:

$$d_{t_1}^* = \sqrt{(8-0)^2 + (1-0)^2} + \sqrt{(8-8)^2 + (-1-1)^2}$$

$$= \sqrt{65} + 2 = 10.06;$$

$$d_{t_2}^* = \sqrt{65} + 2 \cdot \sqrt{2} = 10.89;$$

$$\max_{i \in I_{t_1}}(d_{0,i}) = \sqrt{65} = 8.06;\ \max_{i \in I_{t_2}}(d_{0,i}) = 9;$$

$$\delta_{t_1} = 10.06 - 8.06 = 2;\ \delta_{t_2} = 10.89 - 9 = 1.89.$$

The calculations show that the detour of t_1 exceeds the detour of the tour superset t_2 with $\delta_{t_1} = 2 > \delta_{t_2} = 1.89$. This proves that adding stores to tours can *decrease* the detour. In the context of our problem, this means that tour supersets of an infeasible tour cannot automatically be excluded from the search.

Proposition 1 consequently requires to specifically determine the infeasible supersets of an infeasible tour. Proposition 2 formulates two essential properties of such supersets. □

Proposition 2. *Consider an infeasible tour t of the master problem ($t \in T^{\text{GAP}}$). For any tour superset t' of tour t, the following holds*:

i) *A tour superset t' of an infeasible tour t can only be feasible when including a node j that increases the maximum direct distance $d_{0,i}$ from the starting node 0 on the tour t, with $j \in I, i \in I_t$, and $j \neq i$.*

ii) *If a tour superset t' extending the infeasible tour t by any node j is not feasible, it cannot turn feasible when adding any further node k, with $j, k \in I, i \in I_t$, and $j, k \neq i$.*

Proof of i). According to the triangular inequality, the insertion of a store j into a tour t always leads to a higher or equal tour length. Therefore, the feasibility with regard to the detour of the expanded tour can only be achieved if $d_{0,j} > \max_{i \in I_t}(d_{0,i})$. □

Proof of ii). Let tour t be infeasible as it exceeds the detour limit Δ, and let all supersets t', with $|I_{t'}| = |I_t| + 1$ that extend t by a single node, be infeasible as well. Under these conditions, we prove that any tour superset t'' of t' and t, with $|I_{t''}| > |I_{t'}|$, also exceeds Δ and is infeasible.

Under the triangle inequality, the shortest distance $d^*_{t''}$ for any tour superset t'' as defined above must be greater than or equal to the shortest distance of any superset tour t' of t:

$$d^*_{t''} \geq d^*_{t'}. \tag{35}$$

Further, let node $k \in I$ be any node added to t' such that superset t'' is obtained. We must distinguish two cases:

- Node k does not increase $\max_{i \in I_{t'}}(d_{0,i})$, that is, $d_{0,k} \leq \max_{i \in I_{t'}}(d_{0,i})$. In this case t'' cannot be feasible according to Proposition 2 i).
- Node k increases $\max_{i \in I_{t'}}(d_{0,i})$, that is, $d_{0,k} > \max_{i \in I_{t'}}(d_{0,i})$ and consequently $\max_{i \in I_{t''}}(d_{0,i}) = d_{0,k}$. We know that t' is infeasible for any node added to t such that $\delta_{t'} = d^*_{t'} - \max_{i \in I_{t'}}(d_{0,i}) > \Delta$. This specifically holds for the superset t'_{k} extending t by node k, that is, $I_{t'_{\text{k}}} = I_t \cup \{k\}$. It follows that

$$\delta_{t''} = d^*_{t''} - max_{i \in I_{t''}}(d_{0,i}) = d^*_{t''} - d_{0,k} \geq d^*_{t'_{\text{k}}} - d_{0,k} > \Delta. \tag{36}$$

Equations (36) hold for any $k \in I$, and in particular for node k with $d_{0,k} = \max_{i \in I}(d_{0,i})$.

Considering (35) and (36), the detour of t'' must then always be equal to or larger than the detour of any tour superset t', such that all tour supersets t'' of tour t are infeasible. □

Respecting the properties defined by Propositions 1 and 2, we can formulate the detour cuts with

$$\sum_{i \in I_s} x_{i,t} \leq |I_s| - 1 + \sum_{j \in F_s} x_{j,t} \quad \forall t \in T,\ s \in \{1, \dots, |S|\},\ I_s \in S \tag{37}$$

where S comprises all infeasible assignments found within the subproblems (see Section 3.3.2), and the set F_s contains all stores that repair the infeasible store assignment $I_s \in S$. Equations (37) then state that it is not possible to include all stores from an infeasible assignment I_s in the same tour unless they are combined with a store from the set F_s. The detour cuts (37) forbid the current assignment and all its supersets, apart from supersets including a store from the set F_s.

For all stores $j \in I / I_s, I_s \in S$ with $d_{0,j} > \max_{i \in I_s}(d_{0,i})$, we check all single insertions (see Proposition 2 i)) to I_s, and add stores leading to a feasible tour superset to F_s. This logic leads to the exclusion of the most tours possible. Adding all detour cuts to the master problem turns the GAP into the C-VRP-ZT as all infeasible tours concerning the detour limit are excluded.

(B2) Distance cuts. We further introduce distance cuts as classical no-good cuts. These are a special case of the detour cuts and can only be applied when all supersets of an infeasible tour assignment I_s are infeasible as well, that is, $|F_s| = 0$. To identify these special cases for tours of the master problem, we utilize the fact that the detour restriction can be used to define a global upper bound on the tour length. We define the upper bound for the tour distance $\overline{d}$ as the sum of the maximum distance from node 0 to any i in the problem instance ($max_{i \in I}(d_{0,i})$) and the detour limit Δ: $\overline{d} = max_{i \in I}(d_{0,i}) + \Delta$. For all tour assignments $I_s,\ s \in \{1, \dots, |S|\}$, with $d^*_s > \overline{d}$, we can forbid the tour and all its tour supersets using the cuts of type (38):

$$\sum_{i \in I_s} x_{it} \leq |I_s| - 1 \quad \forall t \in T, s \in \{1, \dots, |S|\}, \quad I_s \in S : d^*_s > \overline{d} \tag{38}$$

Equations (38) ensure that the stores of an infeasible tour cannot be part of any future tours, that is, a tour in the GAP cannot include the complete set of stores I_s (if the shortest distance of I_s is higher than $\overline{d}$). If applicable, the distance cuts exclude a large number of tours since all supersets of an infeasible tour are excluded. Moreover, these cuts strengthen the detour cuts such that the detour cuts are only applied for tour assignments $I_s \in S : d^*_s \leq \overline{d}$.

3.3.4 | Acceleration Techniques

We propose four additional problem-specific acceleration techniques (i)–(iv) to improve the computational performance of our BAC algorithm.

(i) Upper bound of tours and (ii) warm-start procedure. The zone tariffs found in real-world applications usually have specific structures, making it possible to compute an upper bound on the number of tours present in an optimal solution. In general, the tariffs are structured as follows: First, tours covering higher zones are more expensive than tours in lower

ALGORITHM 1 | Complete BAC algorithm.

```
Generate lower bound (o); generate upper bound for tours and warm start solution    ▷ Acceleration (i) and (ii)
Preprocess detour and distance cuts up to cardinality ζ                            ▷ Acceleration (iii)
Calculate distance upper bound d̄
while Improving do
    Solve GAP master problem
    if GAP solution is feasible then
        Pass set of tours T^GAP of GAP solution
        for t ∈ T^GAP do
            Solve O-TSP to determine distance (d*_t) and detour (δ_t)
            if δ_t > Δ and d*_t ≤ d̄ then
                Check for feasible supersets and add corresponding stores to F_s
                Add detour cuts for I_t as lazy constraint to GAP according to Eq. (37)
            end if
            if d*_t > d̄ then
                Add distance cuts for I_t as lazy constraints to GAP according to Eq. (38)
            end if
        end for
        if Any tour t ∈ T^GAP was infeasible then
            Repair the solution                                                     ▷ Acceleration iv)
            if Repaired solution improves the current best solution then
                Pass new feasible solution to the master problem
                Save solution as the current best solution and add optimality cut
            end if
        else
            Save solution as the current best solution and add an optimality cut
        end if
    else
        return last saved solution (optimal solution found)
    end if
end while
```

zones. Second, due to the volume discounts, a lower truck utilization leads to higher costs per load unit. Third, the fixed costs mean that an additional tour introduces further costs. We make use of this tariff structure to compute an upper bound of required tours by solving the MIP model (1–21) for each zone individually. This results in feasible solutions for each zone, where only the corresponding zone-customers are considered. Combining these zone-individual solutions therefore results in a feasible solution of the complete C-VRP-ZT. Given the above-specified properties of the tariff, this solution further provides an upper bound for the total number of tours of the C-VRP-ZT. As a result, there is no solution with more tours that offers lower total costs than the one constructed for the individual zones.

We further warm start our BAC algorithm by providing a feasible starting solution see also [21, 26]. To obtain the start solution, we use the tours resulting from the upper bound calculation, as these represent a feasible solution for the C-VRP-ZT.

(iii) Preprocessing of cuts. The BAC starts without cuts, meaning that the master problem has significantly fewer constraints than the C-VRP-ZT and little information on the tour feasibility concerning the detour limit. To increase the information content, we precompute the feasibility cuts for tours up to a given cardinality ζ and add the cuts to the master problem before starting the solution procedure. In detail, we generate all tours of cardinality ζ and check their feasibility by solving the subproblem for each tour.

(iv) Repairing of tours. Finally, we implement the option to repair GAP assignments that are infeasible in the scope of the C-VRP-ZT see also [27–29]. An infeasible assignment of the master problem might be very close to feasibility (e.g., only one resulting tour is infeasible). We repair infeasible assignments by excluding stores in the order of their index (highest index first) from infeasible tours until the detour limit is fulfilled. If possible, we reinsert each previously excluded store into another tour (first fit, lowest tour index first) while retaining their feasibility. If no tour can accommodate the additional store, we add the store to a new tour. If a repaired solution improves the current best feasible solution, it is passed to the master problem.

3.3.5 | BAC Algorithm

We depict the complete BAC framework in Algorithm 1, showing the interdependencies between the master problem, the subproblem, and the acceleration techniques.

4 | Numerical Results

We begin our numerical studies by introducing the benchmark data set applied (Section 4.1), and then analyze the computational performance of the BAC approach in Section 4.2. We further derive managerial insights related to the tariff structure in Section 4.3. Last, Section 4.4 solves an application case from the retail industry and elaborates on the LSP revenue and retailer cost. All algorithms are implemented in Python (version 3.12) using Gurobi (version 11.0.2) as a solver. We apply Gurobi to solve the master problem and utilize the "MIPSOL" callback function to conduct the feasibility checks within the BAC. We solve the O-TSP subproblem for the feasibility checks using a dynamic programming algorithm [24]. The computations were performed on an AMD 5950X processor (16 cores, 3.4 gigahertz, disabled turbo boost) with 64GB of random access memory. We limit the processor to a single core and apply a time limit of 3600 s.

4.1 | Data Instances

As our work introduces the C-VRP-ZT, there are no common instances available for the problem setting. We create a benchmark data set based on the well-known instances proposed by Solomon [30] to fill this gap. Using this data set allows us to assess the reproducible performance of the BAC and to generalize our findings concerning managerial insights. The data set comprises $3^4 = 81$ instances with the following specifications:

- Three cluster types: We use all geographical distributions of the original Solomon instances for the stores, namely clustered (C), random (R), and random-clustered (RC).
- Three instances sizes: Each Solomon instance includes 100 customers from which we draw 30, 45, and 60 locations to represent small, medium, and large problems.
- Three store settings: We uniformly draw three different sets of store locations per instance size while the DC location is fixed.
- Three demand scenarios: We replace Solomon's demand data with uniformly drawn demands in an interval of [1; 34]. The upper limit of 34 represents the standard truck capacity for Euro pallets. We apply three different demand assignments per instance size.

The further data specification is inspired by the tariff structures found in practice (see Section 2) and the industry application (see Section 4.4). We apply a price for each zone-load combination with a fixed and variable component. The variable cost per loadunit (LU) is subject to volume discounts. The price per zone-load combination is then determined as $\text{price}(\text{zone}, LU) = \text{fixed\_cost}(\text{zone}) + \sum_{i=0}^{LU-1}(\text{base\_cost}(\text{LU}) \cdot (1 - \text{volume\_discounts})^{(i)})$. This formula results in a two-dimensional tariff (see Table 1) with the farthest visited zones as the columns and the discrete vehicle loads as rows. Each column starts with a fixed cost for the vehicle, which increases for higher zones, and adds a variable cost per LU. The variable costs per LU decrease exponentially for a larger number of LU due to the volume discounts. We derive the tariff parameters for our base setting according to direct information from LSPs. The base case considers volume discounts of 5% and sets the absolute detour limit Δ at 10% of the maximum distance from the DC to any store of the instance (i.e., $\Delta = 0.1 \cdot \max_{i \in I}(d_{0,i})$). We further consider five circular zones around the DC with the following structure: The first zone represents a circle with center $c = DC$ and radius r. The second zone then represents the difference between the first zone and a circle with center c and radius $2r$. All further zones follow the same structure. We choose the radius r so that the farthest store from the depot of any Solomon instance is included in the highest zone. Finally, we set the cardinality for the preprocessing of cuts at $\zeta = 2$, as this showed the best performance of the BAC in our pretests. The complete test data is available on https://github.com/NikTuma/CVRPZT.git.

4.2 | Computational Performance

We evaluate the algorithmic performance concerning (a) run time and solution quality, as well as (b) the impact of the acceleration techniques. The specific features of our setting result in a problem structure that hinders direct comparisons with other approaches. Therefore, we apply the following three different methods for the comparison:

- **MIP**: C-VRP-ZT MIP formulation including Equations (1–17) and no acceleration techniques;
- **MIP+**: Accelerated MIP including Equations (1–21) and the acceleration techniques (i)–(iii);
- **BAC**: Proposed algorithm in Section 3.3 incl. all acceleration techniques and valid inequalities.

(a) Run time and solution quality. Table 4 shows that the BAC substantially reduces computation times and solves 65 out of 81 instances (> 80%) to proven optimality within the run-time limit of 3600 s. Only 16 large instances with 60 stores cannot be solved optimally within 3600 s but show a minor average MIP gap of 1.2%. We use the run-time limit for comparability reasons among the solution approaches. Relaxing the limit for the BAC and solving all instances to optimality in a further test results in an average run time of 3763 s. This confirms the BAC's computational efficiency.

The MIP shows the highest run times, solving only 15 instances to optimality within the run-time limit. The MIP+ performs better than the MIP and optimally solves all small and most medium-sized instances (24/27), while none of the large instances could be solved optimally. The last three columns further report BAC statistics showing that the subproblem is consulted sparsely and that the master problem has the main impact on the run time. This effect can be mainly attributed to the preprocessing of cuts. We derive the demand distribution and detour limits from our industry application, which results in a small number of customers per tour—most optimal tours include two stops. Since the BAC algorithm precomputes all cuts for two-stop tours ($\zeta = 2$), only larger tours must be checked for feasibility in the subproblem.

We further test the impact of the regional distribution of stores by analyzing the run time and solution quality differentiated by cluster types C, R, and RC. However, the clustering of stores has only a minor impact, as summarized in Appendix A.

TABLE 4 | Analysis of computational efficiency with Solomon-based instances, average of 27 instances.

	Run time (seconds)[a]			Optimal solutions[b]			MIP gap (%)			BAC SP statistics[c]		
Stores	**MIP**	**MIP+**	**BAC**	**MIP**	**MIP+**	**BAC**	**MIP**	**MIP+**	**BAC**	**Calls**	**Time**	**Cuts**
30	2422	23	7	10	27	27	3.7	0.0	0.0	18	0.01	23
45	3163	1162	147	5	24	27	7.1	0.06	0.0	52	0.04	94
60	3600	3600	3058	0	0	11	> 9.8	> 3.3	0.7	132	0.08	230

[a]Total computation time, incl. master- and subproblem.
[b]Optimal solutions obtained out of 27 instances for each store size.
[c]Subproblem (SP) statistics for the BAC: time in seconds.

TABLE 5 | Average run time if acceleration technique is not applied within BAC in seconds, average of 27 instances; in brackets: Average MIP gap in % at 3600 s.

Stores	BAC	−RT	−WS	−UB	−PC	−VI
30	6.8	7.2	7.9	9.4	374 (0.1)	3600 (15.1)
45	147	149	143	205	3600 (5.2)	3600 (14.1)
60	3058 (0.7)	3075 (0.7)	2972 (0.8)	3487 (1.2)	3600 (8.3)	3600 (13.7)
Avg. increase (factor)		1.03	1.04	1.31	25.95	184.90

Abbreviations: (RT): repairing of infeasible tours; (WS) warm-start solution; (UB) upper bound on the number of tours; (PC) preprocessing of cuts; (VI) valid inequalities.

TABLE 6 | Impact of number of zones, average of 27 instances.

	Run time (seconds)[a]			Cost change			No. of tours			Travel distance[b]		
Zones	**5**	**10**	**20**	**5**	**10**	**20**	**5**	**10**	**20**	**5**	**10**	**20**
30 stores	7	16	24	—	+0.06%	+0.04%	18.9	18.9	18.9	631	632	631
45 stores	147	431	828	—	+0.07%	+0.07%	27.1	27.1	27.1	910	911	913
60 stores	3058	3539	3600	—	+0.08%	+0.05%	35.4	35.4	35.7	1187	1185	1196

Note: Average MIP gaps for 60 store instances: 5 zones: 0.7%; 10 zones: 1.5%; 20 zones: 2.7%.
[a]Run time of BAC.
[b]Total travel distance obtained in post-processing by calculating the sum of O-TSP distances for every tour.

(b) Impact of acceleration techniques. Table 5 highlights the value of the acceleration techniques and the valid inequalities on the average BAC run times across all instances. The column "BAC" describes the run times of the complete algorithm. The further columns report the run times when excluding the respective acceleration. The preprocessing of cuts (PC) and the valid inequalities (VI) have the highest impact. Excluding these accelerations increases run times by a factor of 26 and 185. PC and VI thus ensure that most instances can be solved to optimality within 3600 s and substantially decrease the average MIP gap. The remaining accelerations (RT and UB) show a comparatively moderate effect, while WS slightly increases run time but decreases the MIP gap for 60 stores.

4.3 | Numerical Studies Related to Tariff Structures

The tariff structure defines the retailer's distribution costs and the LSP's margins. The main features of the tariff structure are (a) the number of zones, (b) the volume discounts, and (c) the detour limit. These parameters are defined by the LSP. We additionally analyze how volume discounts impact tours when (d) consolidating orders over periods. This planning of delivery frequency constitutes a main lever for retailers to further save costs.

(a) Impact of number of zones. To test the impact of more granular zoning, we divide each zone of the base case into two and four subzones and adapt the tariff accordingly. This extends the base case from five zones to 10 and 20 zones. In general, a more granular zone design enables a more accurate costing, as it better reflects the actual travel distances. We set identical prices for the first and last zones, independent of the number of zones, to ensure comparability. Table 6 shows that more zones significantly increase the computational effort. However, the total cost barely changes as tours remain almost identical, and the tariff structure implies that the store visited in the highest zone determines the total price. This minor impact on costs and distances suggests that the economic benefit of applying more granular zones is negligible as long as the general pricing of zone distance (i.e., the most remote zone(s)) stays constant. This applies for the LSP as well as for the retailer. Applying five zones in our base setting therefore provides sufficiently granular pricing.

(b) Impact of tariff pricing. Another important feature of the zone tariff is the pricing of zones and loads. While the fixed cost per

TABLE 7 | Impact of volume discounts, compared to base scenario with 5% discount, average of 27 instances.

	Change of total cost				Changes in tours[a]				Tours with > 90% capacity utilization[b]				
Discounts	**0%**	**1.0%**	**2.5%**	**10%**	**0%**	**1.0%**	**2.5%**	**10%**	**0%**	**1.0%**	**2.5%**	**5.0%**	**10%**
30 stores	+62%	+45%	+24%	−27%	43.3%	7.5%	2.7%	4.9%	38.7%	43.8%	44.6%	44.6%	44.4%
45 stores	+66%	+48%	+26%	−28%	40.8%	7.6%	2.6%	3.1%	53.5%	56.8%	57.7%	58.1%	57.6%
60 stores	+69%	+50%	+27%	−29%	43.5%	13.8%	12.0%	10.1%	61.7%	67.4%	67.8%	68.0%	67.6%

[a]Share in % of tours that changed (in vehicle load) compared to 5% discount.
[b]Share of tours with > 90% utilization of truck capacity.

TABLE 8 | Impact of detour limit, average of 27 instances.

	Run time (seconds)[a]			Cost change versus Δ =10%			No. of tours			Travel distance[b]		
Δ	**5%**	**10%**	**20%**	**5%**	**10%**	**20%**	**5%**	**10%**	**20%**	**5%**	**10%**	**20%**
30 stores	4	7	24	+4.3%	—	−4.2%	20.4	18.9	17.4	659	631	630
45 stores	55	147	828	+4.0%	—	−3.4%	29.1	27.1	25.5	947	910	910
60 stores	1131	3058	3600	+2.9%	—	−2.8%	37.3	35.4	33.6	1222	1187	1205

Note: Average MIP gaps for 60 store instances: 5 zones: 0.01%; 10 zones: 0.7%; 20 zones: 2.7%.
[a]Run time of BAC.
[b]Total travel distance obtained in post-processing by calculating the sum of O-TSP distances for every tour.

zone remains constant, the load-dependent variable cost enables volume discounts within the tariff (see Section 4.1).

We analyze the adaptation of the fixed costs in two ways: by changing the absolute level, and by altering the difference between zones. Reducing the absolute level by 50% and 100% primarily affects the total cost, which decreases by 8% and 17%, respectively. It does not affect the total number of tours and only slightly changes the tour configurations. Similar effects can be observed when reducing the difference in fixed costs between zones to 60% and 20% of the base value, decreasing total cost by 4% to 7%.

We further vary the height of the volume discounts and analyze how these changes impact solutions. Table 7 shows that the total cost is very sensitive to the discount rate, as the discount applies to every additional LU. Compared to the base case with a 5% discount, the total cost may increase by up to 69% when no discount is granted and decrease by up to 29% when doubling the discounts to 10%. Interestingly, varying the discount height between 1% and 10% changes only up to one-sixth of the tours. The minor changes can be attributed to capacity and detour limits, which permit only small adjustments. Additionally, the fixed costs per zone remain constant, contributing to the stability of the tours. However, about 40% of the tour configurations change when no discount is applied. Our results highlight that the small difference between no discount and 1% discount leads to different tours and truck loads. The discount rewards higher truck capacity utilization and consolidating orders across various zones. This means giving up lower zone prices for fuller truckloads. As a result, a small discount is already effective to increase the share of full trucks. Counterintuitively, introducing discounts does not lead to fewer tours in this setting. The fixed zone costs in the tariff suffice to minimize the number of tours.

(c) Impact of detour limits. LSPs use the detour limit to control the travel distance on a tour and to force retailers to provide meaningful store combinations. Smaller detour limits restrict the retailer's flexibility when combining stores. Table 8 summarizes results for different absolute detour limits Δ. All values specify a percentage of the maximum distance from the depot to any store of the instance (e.g., medium: $\Delta = 0.1 \cdot \max_{i \in I}(d_{0,i})$). An increasing Δ generally allows the retailer to consolidate more stores on a tour. These additional options drive up computation times but lead to fewer tours (i.e., more customers per tour) and cost savings for the retailer. Increasing the detour limit from 5% to 10% further leads to savings in total travel distances for the LSP. This counterintuitive finding shows that allowing higher detours and thus a more flexible tour building can benefit both parties. Doubling the detour limit again from 10% to 20% does not lead to a further distance decrease. However, the total travel distances still remain below the restrictive 5% scenario. A detour limit of 10% avoids excessive computational complexity and minimizes the travel distance of the LSP, while a higher detour limit enables increased cost savings for the retailer.

(d) Impact of order consolidation. After analyzing the effects of LSP-related parameters, this analysis considers the impact of changes on the retailer side. The retailer may consolidate orders over multiple days to obtain higher delivery volumes per store and benefit from volume discounts. We therefore analyze the effect of adjusting the delivery frequency by aggregating orders over two periods. This enables us to analyze the trade-off between increased order consolidation for smaller orders and order aggregation for increased volume discounts. We first consider a scenario with separate deliveries for two delivery days and corresponding smaller volumes, denoted as *unconsolidated* scenario ("U1", "U2"). The order sizes in this scenario are drawn uniformly for each day, with $q_i \in [1; 17]$. The smaller order sizes

TABLE 9 | Impact of order consolidation over periods, average of 27 instances.

	Demand			Cost change	LSP Revenue/km		No. of tours		Distance[a]	
Stores	U1	U2	C	C versus U1 & U2	U1 & U2	C	U1 & U2	C	U1 & U2	C
30	259	289	548	−5.8%	14.2	17.9	23.7	20.7	918	683
45	396	439	835	−0.9%	14.8	18.4	31.4	30.5	1248	994
60	537	567	1,104	+0.9%	15.0	18.4	39.8	39.7	1583	1298

Note: Average MIP gap for 45 stores: U1: 2.1%; U2: 2.7%; C: 0.0%; Avg. MIP gap for 60 stores: U1: 8.8%; U2: 8.4%; C: 0.1%.
[a]Total travel distance obtained in post-processing by calculating the sum of O-TSP distances for every tour.

TABLE 10 | Application case: Data and results for region West.

Instance data			Demand data				Cost change[a]		No. of tours		Run time
Day	Stores	Zones[b]	Min.	Max.	Avg.	Total	Case	BAC	Case	BAC	BAC
1	34	4	1	26	13.5	458	−14.6%	−22.0%	16	15	62
2	40	4	7	34	15.8	633	−11.9%	−17.7%	24	21	1133
3	25	3	6	30	15.8	396	−7.4%	−16.3%	15	13	18
4	10	2	12	22	17	170	−11.4%	−21.7%	6	5	0.7
5	13	3	7	24	16.2	211	−15.0%	−20.0%	7	7	0.8
Avg.	24.4	3.2	13.7	27.2	13.7	373.6	−12.1%	−19.5%	13.6	12.2	243

[a]Cost change compared to baseline, that is, separate delivery of each store.
[b]Zones: number of tariff zones including at least one store.

potentially allow for higher order consolidation. We then aggregate the orders of these two days to obtain a higher total order volume and denote this as *consolidated* scenario ("C"). Note that we do not allow split deliveries in this scenario.

Table 9 shows the impact of the order aggregation. Looking at the total cost, we see that the consolidation across delivery days mostly benefits retailers with few stores. The lower store density and order aggregation across periods enable the retailer to better exploit the volume discounts and the truck capacities. This reduces the number of tours needed. When the store density increases, this effect decreases (for 45 stores) and even leads to higher total costs for 60 stores. A high store density already implies sufficient consolidation possibilities across the stores to exploit the volume discounts. From the LSP perspective, order consolidation is always beneficial as it reduces the total travel distance in all scenarios. The distance for a separate delivery in U1 and U2 is significantly higher, substantially reducing the LSP's revenues/km compared to the C scenario. In summary, retailers operating a smaller store network can use order consolidation to reduce their transportation costs when a zone tariff applies, while dense store networks benefit from high delivery frequencies to exploit the tariff's volume discounts.

4.4 | Application in Practice

We complement our general insights on the tariff by an industry application. This generates further insights into the use of our model by applying our approach to real-world data.

Data setting of application case. The industry application is based on a collaboration with a European do-it-yourself retailer. The retailer operates two independent delivery areas for distribution in Germany, denoted by the regions West and East. The data contains all stores of the two regions. The stores submit orders one week ahead, and a LSP carries out the daily deliveries from a single DC. We obtained data for a representative planning week consisting of five delivery days, containing the DC and store locations, the daily orders and volumes, the tariff scheme, and the zone layout. The distance matrix includes road distances. The LSP executing the tours for the retailer employs a tariff with five zones and offers a homogeneous fleet with $Q = 34$ LU. The zones have a concentric shape but are clusters of postal code areas, meaning the shape is partly predetermined by the shape of these areas. We report data ranges for the case data as the detailed data and zone tariff cannot be published due to non-disclosure agreements.

Results of real-world application. We compare the cost benefits of the BAC solution with two benchmarks. The first one is denoted as *baseline* and serves as an upper bound as it provides the total cost when all stores are supplied by individual deliveries and no consolidation takes place. The second benchmark represents the solution generated by the retailer and is denoted as *Case*. The retailer's simplified heuristic first sorts the stores by ZIP codes and then assigns nearby stores to the same tour until the vehicle capacity is reached. This is repeated until all stores are assigned. A final check ensures maintaining the detour limit and proposing a feasible solution to the LSP.

Tables 10 and 11 summarize the application case for the two regions. The BAC solves all planning days to optimality in less than 1200 s. This is an important achievement as distribution planning constitutes a daily task. Comparing the BAC and the retailer's solution to the baseline, we see that the BAC shows an average improvement potential of almost 6 percentage point (pps) over the case company's solution for both regions. The

TABLE 11 | Application case: Data and results for region East.

Instance data			Demand data				Cost change[a]		No. of tours		Run time
Day	Stores	Zones[b]	Min.	Max.	Avg.	Total	Case	BAC	Case	BAC	BAC
1	29	2	6	29	14.3	416	−19.5%	−19.5%	14	14	53
2	19	2	8	30	16.1	306	−4.3%	−12.0%	14	11	7
3	31	2	8	31	16.2	502	−8.0%	−15.4%	19	17	21
4	10	2	9	21	13.6	136	−23.0%	−23.0%	4	4	0.6
5	37	3	8	34	15.8	584	−10.0%	−16.4%	20	19	1173
Avg.	25.2	2.2	7.8	29	15.2	388.8	−13.0%	−17.3%	14.2	13	251

[a]Cost change compared to baseline, that is, separate delivery of each store.
[b]Zones: number of tariff zones including at least one store.

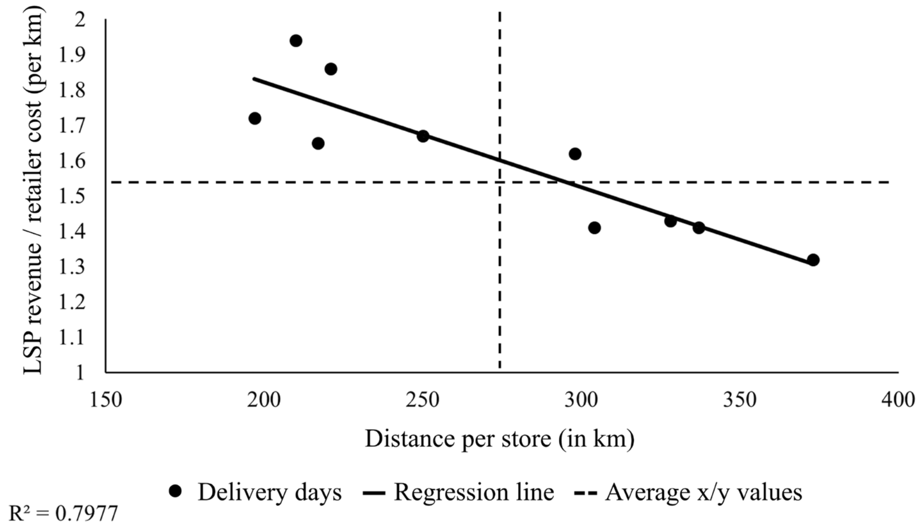

FIGURE 3 | Application case: Revenue and cost factor compared to distance/store across all delivery days.

improvement potential is higher for the region West with an average of 7.4 pps. The BAC results in tours with a higher truck utilization, reducing the total number of required tours from 139 to 126 across regions. The BAC tours for both regions include an average of two stops. This is due to the relatively high average demand per store of 14.5 LUs, equaling almost 50% of the truck capacity. Our findings show that the BAC approach exploits the tariff provided by the LSP, reducing the retailer's cost and the number of tours required. This can further free capacities of LSPs.

Practical implications on cost and revenue. The final analysis compares the solutions from a distance minimization using a Capacitated Vehicle Routing Problem (C-VRP) with those from the C-VRP-ZT. This enables us to analyze the cost and revenue structures of the retailer and the LSP. For these purposes, we adapt the C-VRP-ZT formulation and change the objective to a distance minimization problem, transferring the problem to a C-VRP with open tours and a detour limit, which we solve using Gurobi.

The comparison of the solutions with regard to solution structures shows that both the C-VRP and C-VRP-ZT plan the same total number of tours. However, the C-VRP-ZT results in 2.4% longer travel distances, and 46% of tours differ. Further, we relate the optimal distances of the C-VRP to the C-VRP-ZT cost. The cost of the retailer reflects the revenues for the LSP, and we use both terms interchangeably. Dividing the minimum cost of the C-VRP-ZT by the minimum distances of the C-VRP provides a revenue factor per distance unit for the LSP and a respective cost factor for the retailer. Furthermore, we calculate the distance per store by dividing the travel distance of the C-VRP solution by the number of stores. This refers to the density of the stores that need to be serviced. This comparison enables us to derive insights into the revenue structures of tours built within the C-VRP-ZT, indications about possible margins of LSPs, and a benchmark against the use of a potential own fleet for the retailer.

Figure 3 depicts the relation of the retail cost factor and the distance per store. The analysis helps retailers assess whether it is more cost-effective to use their own fleet or outsource deliveries. Our results indicate that if a retailer's fleet can serve the stores at an average cost of 1.55 € per km or less, using their own fleet would be beneficial. If the retailer's cost is above this threshold, outsourcing to an LSP is more attractive. This, however, also depends on the distance per store. The average distance per store in our case is 274 km. The cost factor amounts to up to 1.94 € per km for delivery days where the store locations lie in closer proximity, whereas it drops to 1.32 € per km for days requiring longer driving distances between stores. This implies for the retailer that insourcing may become more beneficial for denser networks (lower distance/store) and outsourcing for far-distance deliveries. For LSPs, our findings indicate that a shorter distance per store leads to higher revenue per distance unit. We run a linear regression to show the relationship between LSP revenue/retailer cost and the distance per store. The regression further shows that delivery days with a revenue factor above the line are especially profitable for the LSP. We find that a revenue factor above the line is correlated with a high average demand, while a revenue factor below correlates with a low average demand. As a consequence, high-demand delivery days are more beneficial for the LSP, which is also in line with our findings on the impact of consolidation (see above in Section 4.3).

5 | Conclusion

Our work formally studies the real-world transportation problem in which a retailer employs an LSP for deliveries billed using a zone-based tariff. This tariff incorporates a detour limit and volume discounts. The retailer aims to obtain cost-minimal tours while adhering to the LSP requirements. To address the retailer's decision problem, we formalize the problem and introduce a novel mathematical model of the C-VRP-ZT. We enhance the formulation using valid inequalities and develop an exact BAC solution approach. The BAC method decomposes the problem to effectively handle larger instances and employs problem-specific acceleration techniques to boost computational efficiency. We apply benchmark problem instances and show that the BAC algorithm outperforms the standard and accelerated versions of the MIP solved with Gurobi by up to two orders of magnitude. The benchmark instances are publicly available for future comparisons. To round up our study, we solve an application case and derive managerial insights on the structure of the zone tariff and the delivery frequency. Our results show that

- increasing the number of zones in a tariff does not yield meaningful economic benefits for LSPs nor retailers, but rather complicates the problem. Simpler tariff structures with fewer zones are sufficient for accurately pricing logistic services and help reduce complexity;
- minor adjustments to the pricing and discounting structure within the tariff influence both the total cost for retailers and the revenues of LSPs, while the effect on tour configurations is minimal;
- relaxing the detour limits does not necessarily increase travel distances for LSPs. Proper calibration of the detour limit is key to balancing retailer costs, routing distances, and LSP revenues. Allowing flexibility in the detour limit can benefit both parties;
- consolidating orders across delivery days especially benefits retailers operating smaller delivery networks. The consolidation enables bundling effects for them, leading to higher discounts through increased truck loads. In contrast, retailers with a dense delivery network and sufficient order volume may benefit from more frequent deliveries and no additional consolidation;
- comparing the cost of a zone-based tariff to a classic distance-based pricing gives valuable insights on the profitability of using an in-house fleet or employing LSPs. Further, we see that the different objectives (C-VRP vs. C-VRP-ZT) do not result in a different total number of tours and lead to minor differences in total travel distances.

Future Areas of Research Formalizing the decision problem and developing the BAC lays the groundwork for future research in two main areas. From a strategic standpoint, the problem of designing a fair zone tariff can be addressed, and insights on different tariff design options and their impact on the stakeholders can be derived (e.g., different detour limit options and shapes of zones). The problem can be generalized by including time windows, multiple depots, and split deliveries. The model could also be extended to incorporate LSP selection where different LSPs serve customers from one DC or multiple delivery modes (e.g., an own fleet combined with an external fleet). Additionally, the tariff could be generalized to account for further cost drivers. Investigating optimal cooperation models and profit-sharing between the retailer and LSP is also a promising area. Last but not least, the optimization problem of an LSP constitutes an interesting research area. Future research could address the optimal routing and tariff structures of an LSP providing service to multiple retailers.

From an operational standpoint, the BAC can be accelerated by heuristic boosting to derive new feasible solutions that act as optimality cuts. Another promising option for an exact solution algorithm could be a problem-specific implementation of a branch-and-price algorithm, specifically designed to efficiently generate shortest paths with the absolute detour limit in mind. Further, adapting the BAC to a heuristic-based algorithm and reducing computation times at the cost of solution quality can be beneficial when solving large decision problems with prohibitive complexity. We recognize that the novelty of the problem formulation and the absence of established solution approaches present challenges for direct comparative analyses. To support future research and further applications, we provide a comprehensive MIP model, an exact solution approach, and a set of benchmark instances. These contributions lay the groundwork for systematic evaluation, replication, and extension—paving the way for deeper exploration of this relevant industry problem.

Acknowledgments

Open Access funding enabled and organized by Projekt DEAL.

Data Availability Statement

The data that support the findings of this study are openly available in Github at https://github.com/NikTuma/CVRPZT.git.

References

1. Statista, "Third-Party Logistics (3pl): Market Data & Analysis," 2023, https://www.statista.com/study/147130/third-party-logistics-3pl-market-data-and-analysis/.

2. Armstrong & Associates, “3PL Customers Report Identifies Service Trends, 3PL Market Segment Sizes, and Growth Rates,” 2017, https://www.3plogistics.com/3pl-customers-report-identifies-service-trends-3pl-market-segment-sizes-growth-rates/.

3. H. W. Hamacher and A. Schöbel, “Design of Zone Tariff Systems in Public Transportation,” *Operations Research* 52 (2004): 897–908.

4. United Parcel Service, “2024 Ups Rates Daily,” 2024, https://www.ups.com/assets/resources/webcontent/en_US/Daily_US48_Final_CORRECTED.pdf.

5. DHL Express, “Dhl Express Service & Rate Guide 2024 United States,” 2024, https://mydhl.express.dhl/content/dam/downloads/us/en/rate-guide/service_and_rate_guide_us_en_2024.pdf.

6. P. Hansen, D. Peeters, and J. F. Thisse, “Facility Location Under Zone Pricing,” *Journal of Regional Science* 37 (1997): 1–22.

7. F. Azadian and A. Murat, “Service Location Grouping and Pricing in Transportation: Application in Air Cargo,” *European Journal of Operational Research* 267 (2018): 933–943.

8. N. Tuma, M. Ostermeier, and A. Hübner, “Optimal Transportation Planning for a Do-It-Yourself Retailer With a Zone-Based Tariff,” *INFORMS Journal on Applied Analytics* 54 (2024): 312–328.

9. H. Kuhn and M. Sternbeck, “Integrative Retail Logistics: An Exploratory Study,” *Operations Management Research* 6 (2013): 2–18.

10. M. Frank, M. Ostermeier, A. Holzapfel, A. Hübner, and H. Kuhn, “Optimizing Routing and Delivery Patterns With Multi-Compartment Vehicles,” *European Journal of Operational Research* 293 (2021): 495–510.

11. A. Holzapfel, A. Hübner, H. Kuhn, and M. Sternbeck, “Delivery Pattern and Transportation Planning in Grocery Retailing,” *European Journal of Operational Research* 252 (2016): 54–68.

12. G. Soleilhac, F. Lehuédé, J. Medina, and O. Péton, “The Vehicle Routing Problem With FTL and LTL Carriers,” 2022, https://ssrn.com/abstract=4110444.

13. A. Ceselli, G. Righini, and M. Salani, “A Column Generation Algorithm for a Rich Vehicle-Routing Problem,” *Transportation Science* 43 (2009): 56–69.

14. S. Dabia, D. Lai, and D. Vigo, “An Exact Algorithm for a Rich Vehicle Routing Problem With Private Fleet and Common Carrier,” *Transportation Science* 53 (2019): 986–1000.

15. K. A. Lindsey, A. L. Erera, and M. W. Savelsbergh, “A Pickup and Delivery Problem Using Crossdocks and Truckload Lane Rates,” *EURO Journal on Transportation and Logistics* 2 (2013): 5–27.

16. S. Ceschia, L. Di Gaspero, and A. Schaerf, “Tabu Search Techniques for the Heterogeneous Vehicle Routing Problem With Time Windows and Carrier-Dependent Costs,” *Journal of Scheduling* 14 (2011): 601–615.

17. R. D. Tordecilla, P. J. Copado-Méndez, J. Panadero, L. d. C. Martins, and A. A. Juan, “An Agile and Reactive Biased-Randomized Heuristic for an Agri-Food Rich Vehicle Routing Problem,” *Transportation Research Procedia* 58 (2021): 385–392.

18. E. S. Thorsteinsson, “Branch-and-Check: A Hybrid Framework Integrating Mixed Integer Programming and Constraint Logic Programming,” in *Principles and Practice of Constraint Programming—CP 2001: 7th International Conference, CP 2001 Paphos, Cyprus, November 26–December 1, 2001 Proceedings* (Springer, 2001), 16–30.

19. M. L. Fisher and R. Jaikumar, “A Generalized Assignment Heuristic for Vehicle Routing,” *Networks* 11 (1981): 109–124.

20. C. Kloimüllner and G. R. Raidl, “Full-Load Route Planning for Balancing Bike Sharing Systems by Logic-Based Benders Decomposition,” *Networks* 69 (2017): 270–289.

21. R. F. Fachini and V. A. Armentano, “Logic-Based Benders Decomposition for the Heterogeneous Fixed Fleet Vehicle Routing Problem With Time Windows,” *Computers & Industrial Engineering* 148 (2020): 106641.

22. J. Benders, “Partitioning Procedures for Solving Mixed-Variables Programming Problems,” *Numerische Mathematik* 4 (1962): 238–252.

23. R. Rahmaniani, T. G. Crainic, M. Gendreau, and W. Rei, “The Benders Decomposition Algorithm: A Literature Review,” *European Journal of Operational Research* 259 (2017): 801–817.

24. M. Held and R. M. Karp, “A Dynamic Programming Approach to Sequencing Problems,” *Journal of the Society for Industrial and Applied Mathematics* 10 (1962): 196–210.

25. G. Codato and M. Fischetti, “Combinatorial Benders' Cuts for Mixed-Integer Linear Programming,” *Operations Research* 54 (2006): 756–766.

26. M. Kang and C. Lee, “An Exact Algorithm for Heterogeneous Drone-Truck Routing Problem,” *Transportation Science* 55 (2021): 1088–1112.

27. A. Ciré and J. N. Hooker, “A Heuristic Logic-Based Benders Method for the Home Health Care Problem,” in *Proceedings of the 2012 Matheuristics* (Fourth International Workshop on Model-Based Metaheuristics, 2012).

28. M. Riedler and G. Raidl, “Solving a Selective Dial-a-Ride Problem With Logic-Based Benders Decomposition,” *Computers & Operations Research* 96 (2018): 30–54.

29. P. Fontaine and S. Minner, “A Branch-And-Repair Method for Three-Dimensional Bin Selection and Packing in e-Commerce,” *Operations Research* 71 (2023): 273–288.

30. M. M. Solomon, “Algorithms for the Vehicle Routing and Scheduling Problems With Time Window Constraints,” *Operations Research* 35 (1987): 254–265.

Appendix A

Impact of Store Distribution

Table A1 shows the impact of the geographical distribution of the stores within the Solomon-based data (C, R, RC). The BAC provides the best run times and highest solution quality across all instances. The geographical store distribution marginally affects the run time of the BAC. The algorithm is especially efficient on the RC instances.

TABLE A1 | Analysis of Computational Performance With Different Geographical Distribution of Stores, Average of Nine Instances.

		Run time (seconds)			Optimal solutions[a]			MIP gap (%)		
Stores	**Approach**	**C**	**R**	**RC**	**C**	**R**	**RC**	**C**	**R**	**RC**
30	MIP	2296	2493	2476	4	3	3	3.4	2.7	5.0
	MIP+	43	19	6	9	9	9	0.0	0.0	0.0
	BAC	7	10	4	9	9	9	0.0	0.0	0.0
45	MIP	3015	3600	2874	2	0	3	4.9	10.4	6.0
	MIP+	1749	1348	389	7	8	9	0.1	0.1	0.0
	BAC	184	143	116	9	9	9	0.0	0.0	0.0
60	MIP	3600	3600	3600	0	0	0	9.5	9.6	> 10.3
	MIP+	3600	3600	3600	0	0	0	3.5	3.8	2.6
	BAC	3057	3216	2900	4	2	5	0.7	0.8	0.6

[a]Number of optimal solutions obtained out of nine instances for each store size.