Received: 30 November 2019 | Revised: 4 September 2020 | Accepted: 10 September 2020 | Published on: 10 November 2020

DOI: 10.1002/nav.21954

RESEARCH ARTICLE

WILEY

Same-day deliveries in omnichannel retail: Integrated order picking and vehicle routing with vehicle-site dependencies

Daniel Schubert[1] | **Heinrich Kuhn**[1] | **Andreas Holzapfel**[2]

[1]Supply Chain Management & Operations, Catholic University of Eichstätt-Ingolstadt, 85049 Ingolstadt, Germany
[2]Department of Logistics Management, Hochschule Geisenheim University, 65366 Geisenheim, Germany

Correspondence
Heinrich Kuhn, Supply Chain Management & Operations, Catholic University of Eichstätt-Ingolstadt, 85049 Ingolstadt, Germany.
Email: heinrich.kuhn@ku.de

Funding information
Erich-Kellerhals-Stiftung (Foundation), 85049 Ingolstadt, Germany.

Abstract

Currently, the most common delivery mode offered in online retailing is next-day delivery. Retailers, however, are investigating whether same-day deliveries can be offered to online and store customers in order to meet increasing customer requirements. The present paper considers an integrated order picking and vehicle routing problem assuming same-day delivery in the field of omnichannel retailing. Within this context the paper focuses on the picking and delivery process for large durable consumer goods, for example, consumer electronics and home appliances with the option of associated installation services at the customer's home as well as store deliveries for click-and-collect customers and the short-term replenishment of store inventory. We develop a decision support model and a general variable neighborhood search-based algorithm that tackle the problem described by considering the tradeoff between picking and delivery costs while ensuring customers' delivery time windows. The results show the benefits of an integrated solution approach amounting to approx. 13% total cost savings on average, compared to a sequential approach typically applied in retail practice. A case study from an internationally operating omnichannel retailer demonstrates the advantages and applicability of the modeling and solution approach suggested.

KEYWORDS

general variable neighborhood search, online retailing, order picking, retail operations, same-day delivery, vehicle routing

1 | INTRODUCTION

The increasing demand for multi- and omnichannel shopping opportunities is forcing retail companies to continuously enhance their supply chain strategies to satisfy customer requirements (Agatz, Fleischmann, & van Nunen, 2008; Gallino & Moreno, 2019). The online channel is growing especially rapidly. While total revenues of the online channel in Germany amounted to EUR 20.2 billion in 2010, they dramatically increased to EUR 53.3 billion and EUR 57.8 billion in 2018 and 2019, respectively (HDE, 2019). Revenues from household appliances, for example, increased from EUR 2.44 billion in 2015 to EUR 4.30 billion in 2018, equating to a growth of 76.2%. Similarly, revenues increased by 56.0% in the field of electronics and telecommunication from EUR 7.55 billion (2015) to EUR 11.79 billion (2018) (bevh, 2019). Retailers have to face up to this development and to meet ever growing customer expectations, very short delivery times being one of the toughest (Hübner, Holzapfel, & Kuhn, 2015, 2016; Hübner, Holzapfel, Kuhn, & Obermair, 2019). Currently, the standard mode of delivery is "next-day delivery," for example, customer orders that are placed by 8 PM are expected to be fulfilled on the following day (Wollenburg, Holzapfel, Hübner, & Kuhn, 2018). In the meantime, however, retailers are investigating whether and under what conditions same-day deliveries can be offered to customers in order to cope with emerging customer requirements for shorter delivery times, enhancing customer services as a result (McKinsey, 2017; Voccia, Campbell, & Thomas, 2017).

Online retailers typically operate distribution centers (DCs) that are responsible for fulfilling most customer orders (Holzapfel, Kuhn, & Sternbeck, 2018). The products relating to these orders are picked in the DC and subsequently delivered to the locations requested by the customers. In retail practice, both processes, that is, picking and delivery, are mostly considered independent operations (Kuhn, Schubert, & Holzapfel, 2020; Moons, Ramaekers, Caris, & Arda, 2018; Schubert, Scholz, & Wäscher, 2018). Both processes, however, are strongly interconnected due to the short time span between the arrival of customer's order and the delivery time-slot in the context of same-day delivery. This situation applies both where retailers have their own delivery fleet and when they charter a forwarding agent to fulfill their deliveries.

A vehicle, for example, can only start executing its tour once all items from the customer orders assigned have been picked and placed at the loading dock of the warehouse as well as subsequently loaded into the shipping space of the vehicle. The time required to prepare all a tour's orders at the loading dock, however, depends on the order picking schedule. Likewise, the tour's entire delivery time depends on the orders assigned. Thus, a schedule for order picking operations may be favorable for picking but unfavorable for delivery and vice versa. The question therefore arises, which of all possible combinations will minimize overall picking and delivery costs while ensuring agreed delivery dates and service requirements. This is especially crucial for same-day deliveries since additional picking shifts and delivery tours are required during the late afternoon and evening. This increases delivery costs and requires higher flexibility within the retailer's entire distribution system. Retailers therefore hesitate to offer this service, largely because of its unprofitability (Hübner et al., 2016).

Retail practice requires novel decision support systems that tackle the problem described by considering the tradeoff between picking and delivery costs while ensuring customers' delivery time windows. The literature discusses this class of decision problems as integrated order picking and vehicle routing problems (Kuhn et al., 2020; Moons et al., 2018; Schubert et al., 2018). This kind of problem class, however, is difficult to solve since the individual problems are already known to be NP-hard. Exact solution methods are only suitable for very small instances. A retailer, however, has to manage a considerable number of orders even in the case of same-day delivery: up to 200 orders are expected by the company in our case. This means that appropriate heuristic solution procedures are required.

The present paper contributes to the literature in the following way. It takes up the decision problem described and develops a novel mixed-integer program (MIP) that also considers additional installation services at delivery destinations. The paper develops a variable neighborhood search (VNS) based approach that makes it possible to obtain high-quality solutions within short computing times for instances of practically relevant sizes. The solution procedure distinctly improves the results of sequential approaches usually applied in practice. A joint case study with a leading European electronics and household appliance retailer demonstrates the applicability of the modeling and solution approach developed.

The remainder of the paper is organized as follows. Section 2 details the context of the picking and delivery process under consideration and illustrates the interdependencies between both processes. Section 3 reviews related literature, while Section 4 formulates the decision support model of the integrated planning approach. An enhanced VNS, that is, a general VNS (GVNS) approach for solving this problem, is proposed in Section 5. Section 6 comprises numerical experiments with simulated problem instances as well as instances from practice. The same section also presents several managerial insights. Section 7 completes the paper and outlines further research opportunities.

2 | PROBLEM DESCRIPTION

The problem under investigation considers an omnichannel retailer operating a large number of bricks-and-mortar stores and an online channel (Gallino & Moreno, 2019; Hübner et al., 2015; Wollenburg, Hübner, Kuhn, & Trautrims, 2018). The retailer offers large durable consumer goods, for example, consumer electronics and home appliances (TVs, satellite dishes, washing machines, refrigerators, dishwashers, etc.) in both channels. Online as well as store customers can ask for a home delivery of the goods purchased and additionally order associated installation services at the delivery destination requested. In addition, the given structure enables the retailer to fulfill click-and-collect orders for customers who wish to pick up the goods they have ordered at the retailer's stores, or to provide general store deliveries at short notice (Paul, Agatz, Spliet, & Koster, 2019). The retailer fulfills all these different kinds of customer and store orders from one of its omnichannel regional warehouses. This general structure also enables the retailer to provide deliveries to both customers and store locations within a few hours, which is especially important when offering a same-day delivery option (McKinsey, 2017; Voccia et al., 2017). The retailer in our case is able to offer a same-day delivery option for a large number of customer and store orders.

In the remainder of this section we define the timeline of events, the individual fulfillment processes, the resources required and the interconnection between these processes, which requires an integrated planning approach in the same-day delivery setting described. Please note that in the following we denote all orders as "customer orders" regardless of whether they originate from pure online customers,

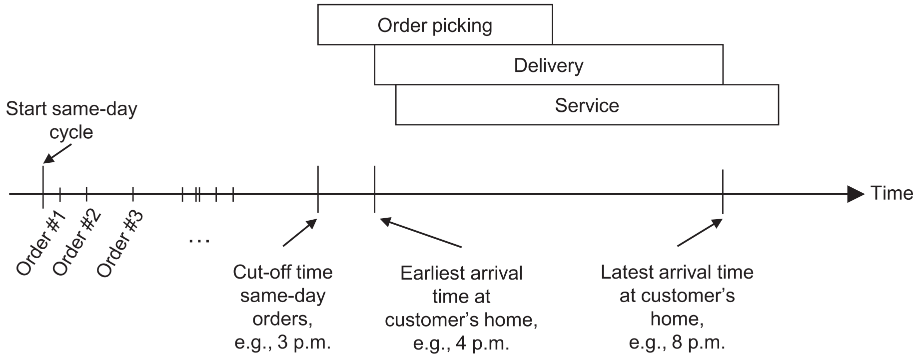

FIGURE 1 Timeline of events and individual processes

click-and-collect customers, store customers, or from the store itself.

2.1 | Timeline of events

In the same-day delivery context, all sorts of customers can place their orders at an early stage of the day, that is, by a certain cut-off time (e.g., 3 PM), and receive the corresponding goods at home in the afternoon or evening. The retailer defines an earliest (e.g., 4 PM) and a latest arrival time (e.g., 8 PM) at the customer's home. The picking process can immediately start after the cut-off time. Note that it is indeed common that same-day delivery orders are not picked beforehand, that is, as soon as these orders arrive, because of space limitations at the loading docks of the warehouse, and to avoid mixing up these orders with standard delivery orders (see, e.g., Hübner et al. (2015)). The picking process of these orders therefore starts after the cut-off time. The delivery process then starts as soon as the respective vehicles are loaded, which could be considerably in advance of the earliest arrival time defined because of the travel time to the customer. The service process, however, can last longer than the latest permissible arrival time at the customer's home because of the service time required. Figure 1 illustrates the timeline of events and the individual processes.

2.2 | Order picking process

Large items are usually picked one by one from the storage area of the warehouse using a special picking device, for example, a fork lift with clamp. Note that this picking strategy is generally denoted "*discrete picking*." A discretely picked item often corresponds to the customer order size since large durable consumer goods are generally individually ordered. The picking process then has the following sequence: an order picker receives a picking order at the dispatching area of the warehouse, drives or walks to the corresponding storage location, collects the item, and finally returns to the unloading area, that is, the ramp of the warehouse, or loads the item directly into the cargo space of the waiting vehicle (De Koster, Le-Duc, & Roodbergen, 2007). The expected processing time of a customer order in the warehouse (pick time) can therefore be approximated beforehand (Schubert et al., 2018). The pick time of a customer order includes a setup time for preparing a (picking) tour, the travel time, a search time for identifying an item, a pick time for the retrieval process, and the loading time into the vehicle (Tompkins, White, Bozer, & Tanchoco, 2010). In addition, the following assumptions must be ensured in order to approximate the expected process times with sufficient accuracy: a permanent application of a specified routing strategy, a warehouse layout with wide aisles, and continuous availability of the items in demand.

2.3 | Delivery process

The delivery process starts as soon as a vehicle is loaded with the products required for the delivery tour scheduled. The earliest possible arrival time at customers' homes has to be considered within this context. In general, two delivery time options are applied in retail practice: (a) customers can choose a certain time window (e.g., 2 hours) within which the corresponding service should start, or (b) the delivery service may start at any time during the entire delivery time span predefined by the retailer, that is, between the earliest and latest arrival time at customer's home. The latter option is denoted "general time window." Furthermore, a delivery tour is typically limited by time and space constraints. The delivery time includes the travel, waiting and service times required to fulfill all customer orders assigned to a tour. In addition, the loading space of the vehicle has to be considered. A typical tour contains 5 to 15 customer orders in the delivery context considered.

2.4 | Additional delivery services

The retailer offers additional installation services at the delivery destination that can be ordered by the customers at an extra charge. In general, three service types are offered in the case of large durable consumer goods: (a) delivery only to the place of destination, (b) delivery to the place of destination and basic installation services, and (c) delivery to the place of destination with all types of installation, for example, including high voltage current connections and/or carpentry services. A qualified employee (driver and/or co-driver) is

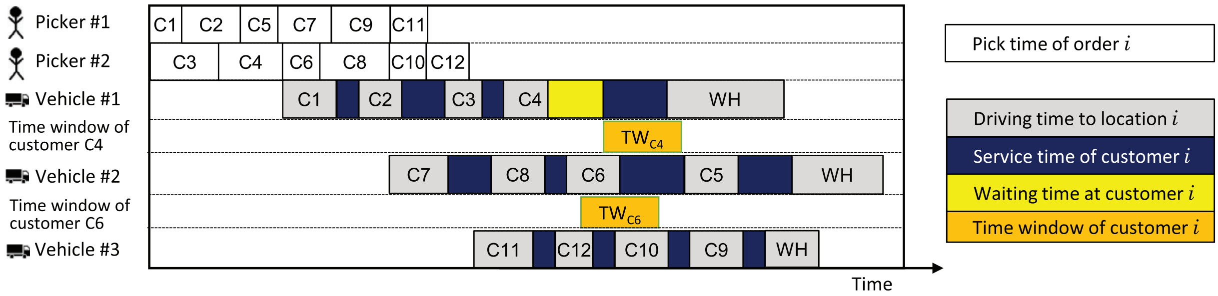

FIGURE 2 Visualization of the interdependencies between picking and delivery process

necessary for each service type, and installation equipment also has to be prepared in the vehicle used. The equipment required and qualifications of the drivers for each service type differ significantly. While there are only a few requirements for the most basic service level, particular equipment is necessary for the highest service level, for example, consumable materials, tools, as well as a specific qualification level of the driver and/or co-driver. The highest service level required by at least one customer order assigned to a van defines the associated vehicle type. This, in addition, defines the minimum qualification level of the driver and/or co-driver. The specialized materials and tools required to fulfill the respective services are either already available in all vans or present in boxes or roll cages that are additionally loaded into the van upon departure. Vehicles can therefore be divided into different vehicle types according to the service types they can be used for and the associated driver qualifications. Each service type requires a certain service time that has to be considered when scheduling the delivery of customer orders.

2.5 | Resources required

Customer demand varies from day to day in retail practice. The resources required will therefore also vary from day to day, that is, the number of vehicles per type and the number of order pickers necessary to process the respective customer orders. In the present paper we assume that these resources are provided by delivery companies, subcontractors performing the picking process, or other business units of the retailer at short notice. Please note that in the present case the delivery company and the subcontractor are willing to offer this short-term flexibility since all parties involved are interlinked within a broad, nationwide and long-standing business relationship. The use of these resources is then remunerated or compensated based on their respective use. The retailer, however, has the entire planning authority and decides at short notice how many units of each of the resources are required to fulfill all customer orders within the planning horizon. The aim of the retailer is therefore to minimize the total costs for pickers and vehicles used while fulfilling all customer orders, including the respective services requested. This type of organization is common practice in the planning environment being considered.

2.6 | Decisions and their interdependencies

In the context of same-day delivery, the order picking process in the warehouse and the home-delivery process, that is, the delivery tours, are interconnected because of the short time span between the arrival of the customer order and the customer's delivery date. The connection of both subproblems is illustrated in Figure 2.

Two pickers and three vehicles are planned to pick and deliver the 12 pending customer orders, that is, C1, C2, …, C12. All orders are allowed to be delivered during the general delivery period. Only customers C4 and C6 have chosen a specific time window for their delivery. A vehicle can start its associated delivery tour as soon as the warehouse has prepared the last customer order assigned to it. A delivery tour generally alternates between travel and service processes. An additional waiting time occurs if a vehicle arrives at the customer's home before the earliest possible arrival time. Vehicle #1, for example, arrives at customer C4 before the associated delivery is requested. The driver and/or service person therefore have to wait until the corresponding time window is reached. The tour for the second vehicle (#2), on the other hand, is scheduled such that it arrives exactly within the time window requested by customer C6.

The integrated planning of both processes, that is, picking and delivery, makes it possible to start the delivery tours as early as possible, which allows more customers to be served within the respective tours. Vehicles are therefore used most efficiently with respect to loading space and driving time. This would not be the case if all delivery tours started at the same time, as is assumed in sequential planning approaches (Klapp, Erera, & Toriello, 2018). The order picking process in the warehouse and the home-delivery process should therefore be coordinated.

In addition, the number of pickers assigned influences the earliest starting time of the respective delivery tours and therefore influences the number of vehicles required. Delivery tours can start earlier with a larger number of pickers, which may reduce the number of vehicles required. On the other hand, a greater number of vehicles allows operation of the warehouse with a smaller number of pickers. This means more time is available to make the products available at the ramp of the warehouse.

2.7 | Decision problem

The present planning problem necessitates determining both the number of pickers and the number of vehicles required per vehicle type. In addition, the customer orders available have to be assigned to the individual pickers and vehicles, and the respective sequence of their execution has to be determined such that the customer time windows promised and the service type required, that is, standard, medium, or full service, are satisfied. The aim of the present paper is therefore to develop a decision support model and an appropriate solution approach that tackles the problem described by considering the tradeoff between picking and delivery costs while ensuring customers' delivery time windows.

3 | LITERATURE REVIEW

The literature review first discusses the literature related to the individual order picking (see Section 3.1) and vehicle routing (see Section 3.2) subproblems. Afterwards we analyze the literature that considers a simultaneous planning approach (see Section 3.3).

3.1 | Order picking problem

On the operational planning level, the number of order pickers per day or shift have to be determined and afterwards assigned to the different warehouse areas and zones. The corresponding planning problems are denoted as workforce level and assignment problems (van Gils, Ramaekers, Caris, & Cools, 2017). The order assignment problem then defines how customer (or picking) orders are assigned to an individual picker (Scholz, Schubert, & Wäscher, 2017). Two main strategies can be distinguished: individual order picking, that is, "discrete picking," and batch picking (De Koster et al., 2007; van Gils, Caris, Ramaekers, & Braekers, 2019; van Gils, Ramaekers, Caris, & de Koster, 2018).

The circumstances of the planning problem considered in the present paper require that customer orders are picked discretely, that is, one by one. The discrete picking problem has similarities with the parallel machine scheduling problem by interpreting order pickers as machines, customer orders as jobs, and pick times as machine-independent processing times (Schubert et al., 2018). Order-specific deadlines also have to be considered in the event that delivery tours are given. A general deadline has to be defined if order picking and vehicle routing are planned independently of each other. All items have to be made available at the ramp of the warehouse at this deadline. In that case the order picking problem is equivalent to the well-known NP-complete bin packing problem (Morihara, Ibaraki, & Hasegawa, 1982). An order picker is interpreted as a bin and its capacity corresponds to the predefined general deadline. The weights of the items equal the pick times of the orders, which have to be assigned to the bins, that is, the pickers. The number of pickers (bins) required to perform all orders is then minimized. A recent overview of the classic (one-dimensional) bin packing problem can be found in Delorme, Iori, and Martello (2016), including both exact and heuristic solution approaches. Morihara et al. (1982) propose three approximation algorithms and compare them to the largest processing time rule (LPTR). LPTR arranges orders according to non-increasing processing times. Applying this sequence, they are successively assigned to the bins with the largest remaining capacity. Morihara et al. (1982) figure out that LPTR is a very competitive approach for minimizing the number of bins (pickers) required.

De Koster et al. (2007) and van Gils et al. (2018) present comprehensive overviews on diverse order picking planning problems, including routing and workforce allocation problems.

3.2 | Vehicle routing subproblem

A classical vehicle routing problem (VRP) basically involves two types of decision: customer orders have to be assigned to vehicles and each vehicle has to be scheduled such that all customer orders are fulfilled. It has been extended over recent decades via additional constraints and decisions. These have most commonly been capacity and time window constraints (Vidal, Crainic, Gendreau, & Prins, 2013a). Three main features particularly characterize the vehicle routing subproblem being considered. First, the service at the customer's home has to be initiated within a specific time window. Second, each customer order contains a specific release date. Third, a customer order can only be served by a subset of the vehicle fleet. The subsequent review therefore concentrates on the literature considering VRPs that include at least one of these features.

3.2.1 | VRPs with time windows

Recent and comprehensive overviews on VRPs that include VRPs considering time windows can be found in Vidal et al. (2013a) and Toth and Vigo (2014). Vidal, Crainic, Gendreau, and Prins (2013b) present a state-of-the-art algorithm for a broad class of vehicle routing problems assuming customer-specific time windows. However, release date constraints are not included in the approaches described and analyzed in these studies.

3.2.2 | VRPs with release dates

The solution to the order picking subproblem defines completion dates of the picking orders, which in turn define order release dates of the delivery orders. The vehicle routing subproblem has to consider these release dates. Archetti, Feillet, and Speranza (2015) investigate the complexity of routing problems with release dates. A multi-period VRP with release and due dates is studied and solved by Archetti, Jabali, and Speranza (2015). Cattaruzza, Absi, and Feillet (2016) present a hybrid genetic algorithm for the multitrip vehicle routing problem with time windows and release dates.

3.2.3 | VRPs with site dependencies

In the present study vehicles are classified into different vehicle types according to the service type requested by customers. VRPs with site dependencies take this into account. Cordeau and Laporte (2001) show that the site-dependent vehicle routing problem (SDVRP) can be formulated as a special case of the periodic vehicle routing problem (PVRP) where the number of delivery periods equals the number of vehicle types, and each customer has to be served on one and only one day. The set of allowable days of visit for each customer then equals the set of vehicles of the corresponding SDVRP that are permitted. This is also true for the multidepot vehicle routing problem (Cordeau, Gendreau, & Laporte, 1997). In that case each depot represents a certain type of vehicle.

We refer to Vidal et al. (2013b) for an extensive review of all of these classes of vehicle routing problems. The current state-of-the-art algorithm for solving the VRP with site dependencies and hard time windows is provided by Vidal et al. (2013b). The authors suggest a hybrid genetic algorithm with advanced diversity control. Other authors suggest neighborhood-based solution approaches, for example, a unified tabu search by Cordeau, Laporte, and Mercier (2004), an adaptive large neighborhood search by Pisinger and Ropke (2007), and a parallelized unified tabu search by Cordeau and Maischberger (2012). Zare-Reisabadi and Mirmohammadi (2015) propose an ant colony system-based approach solving a VRP with site dependencies assuming soft time windows.

3.3 | Integrated and related problems

Only a few publications suggest an integrated approach for solving an order picking and vehicle routing problem. Schmid, Doerner, and Laporte (2013) introduce the integration of order picking and vehicle routing as a worthwhile endeavor for future research. Schubert et al. (2018) consider a problem that integrates order picking and vehicle routing on an operational level. The authors assume a predefined number of order pickers, a given number of homogeneous vehicles and customer-specific delivery dates. In order to minimize total tardiness, that is, the sum of all non-negative differences between scheduled and requested delivery dates, a model formulation and an iterated local search algorithm are proposed. The problem formulated by Moons et al. (2018) concerns a given number of order pickers that can be increased up to a certain level. For delivery operations, a restricted number of vehicles with heterogeneous capacities and costs are assumed to be available at the warehouse. A model formulation is presented that aims to minimize total costs composed of variable costs for the proportional wage of regular and additional order pickers, respectively, fixed costs per vehicle used, and variable costs per travel time of a vehicle. Moons, Braekers, Ramaekers, Caris, and Arda (2019) and Ramaekers, Caris, Moons, and van Gils (2018) deal with an identical problem setting. Moons et al. (2019) develop a heuristic solution approach and Ramaekers et al. (2018) apply the heuristic developed to gain managerial insights regarding time windows. Kuhn et al. (2020) integrate order batching, order picking and delivery operations into a simultaneous planning approach when supplying micro stores and gas station convenience stores with perishable and durable goods. They propose an extension of an adaptive large neighborhood search (ALNS) metaheuristic denoted as general ALNS (GALNS) and demonstrate the applicability of the modeling and solution approach suggested in retail practice.

In contrast to the integrated approaches of order picking and vehicle routing just mentioned, there are several contributions that consider the integrated production and distribution problem, implying a make-to-order production strategy. Chen (2010) and Moons, Ramaekers, Caris, and Arda (2017) provide extensive reviews on this topic. Schubert et al. (2018) also extensively review these kinds of contribution focusing on integrated machine scheduling and vehicle routing problems. However, these approaches are of minor interest in our setting since the present paper focuses on the integrated order picking and vehicle routing problem considering site-dependent constraints, that is, compatibility dependencies between customer sites and vehicle types. To the best of our knowledge, a site-dependent vehicle routing problem (SDVRP)—especially the nested structure of the utility of different vehicle types to fulfill certain types of services—has not been considered part of an integrated problem with any machine scheduling or order picking configuration.

Summarizing the literature available on integrated order picking and vehicle routing problems, only a very limited number of papers tackle the general problem (Kuhn et al., 2020; Moons et al., 2018; Schubert et al., 2018), and to the best of our knowledge the planning problem at hand has not been considered in the literature so far. The present paper therefore contributes to the literature as follows. A novel MIP is developed that tackles the integrated order picking and vehicle routing problem defined considering the tradeoff between picking and delivery costs while ensuring customers' delivery time windows and the individual services requested. The model especially considers additional services at the delivery destination. These individual services require specific qualifications of the delivery staff and a certain vehicle type. These services are nested-organized such that higher qualified teams can fulfill services requiring lower qualifications, but not vice versa. An appropriate heuristic solution procedure is developed that improves the results of the sequential approaches usually applied in practice. In addition, the solution approach developed solves instances of practically relevant sizes within a short computational time. Within an extensive numerical analysis the paper demonstrates the value of integration for order picking and delivery planning, which becomes especially true for omnichannel retailers offering a same-day delivery option. A case study is also conducted in cooperation with a leading European omnichannel retailer that demonstrates the real-life applicability of the modeling and solution approach suggested.

4 | MATHEMATICAL MODEL

In the present section we develop a decision support model for the integrated order picking and vehicle routing problem being considered. To achieve this we first define the parameters specifying the customer orders given at the beginning of the planning horizon. We then name the parameters and decision variables defining both subproblems, the order picking subproblem, denoted **O**rder **A**ssignment and **S**equencing **P**roblem (OASP), and the delivery subproblem, named **V**ehicle **R**outing **P**roblem with **S**ite **D**ependencies (VRPSD). Finally, we formulate an integrated MIP, which is denoted **OAS-VRPSD**.

4.1 | Customer orders

At the time of planning, a set of customer orders $N = \{1, 2, ..., i, ..., |N|\}$ is given and all orders have to be fulfilled on the respective day. Each customer order i is characterized by an item with a certain weight w_i and space requirement v_i, $i \in N$. These items have to be prepared in the warehouse and delivered to the corresponding customer location. Moreover, a customer demands a certain but single type of service that can be fulfilled by specific vehicles of type k, $k \in K$, that include specific service teams that are qualified to fulfill the service required. The binary parameter e_{ik} then indicates whether customer i can be served by vehicle type k ($e_{ik} = 1$) or not ($e_{ik} = 0$). Please note that it is assumed that a higher qualified service team, that is, vehicle type, can also perform lower-level services. Fulfilling the service requested by customer i requires a specific service time s_i, $i \in N$. The length of service time only depends on the type of service to be provided, and is independent of the qualification level of the team providing the service. The service time includes unloading operations and the time for connection and installation according to the service type requested. The service has to start within a customer-specific time window and has to be completed within only one visit. The time window is defined by a lower bound α_i and an upper bound β_i, $i \in N$. The start of service is postponed until the lower bound of the time window is reached if a vehicle arrives early at a customer's home. Late arrivals, that is, a start of service after the upper time limit, however, are not allowed. Consequently, the service requested by customer i has to start exactly within the defined time window $[\alpha_i, \beta_i]$. The starting time of service i is denoted as a_i. Note that this defines the most general case for the problem being considered since general time windows and a delivery deadline are special cases of customer-specific time windows.

4.2 | Order picking subproblem

The item ordered by customer i is customer-specifically collected by a picker on a so-called picking tour from its storage location in the warehouse. The pick time required for collecting the item and loading it into the shipping space of the vehicle is denoted p_i, $i \in N$. We assume the availability of a sufficiently large number of pickers, whereby each of the pickers used is charged at a fixed cost rate of c_{picker}. Determining the number of pickers used is part of the decision-making process. Moreover, customer orders have to be assigned to the pickers and the orders assigned have to be sequenced. The assignment and sequencing decisions are quantified by the following decision variables: $y_i \in \{0, 1\}$ defines whether order i is assigned ($y_i = 1$) to the first position on the picking list of one of the pickers available or not ($y_i = 0$). The sum $\sum_{i \in N} y_i$ then quantifies the number of pickers used. An artificial customer order $|N| + 1$, $N^+ = N \cup \{|N| + 1\}$ and the binary variable z_{ij}, $i \in N$, $j \in N^+$, are introduced to represent the subsequent orders on each of the picking lists (Biskup, Herr-mann, & Gupta, 2008). Variable z_{ij} then indicates whether an order i is processed immediately before order j ($z_{ij} = 1$), or not ($z_{ij} = 0$). Order i is processed last on a picking list if $z_{i,|N|+1} = 1$. This means the point in time when the item of an order is prepared by the warehouse for delivery, that is, the release date r_i, $i \in N$, of the respective order depends on the orders previously assigned to the same picker, that is, the sum of their pick times, and on the pick time of the corresponding order itself.

4.3 | Vehicle routing subproblem

The items prepared by the pickers at the ramp of the warehouse are delivered to the customers by vehicles that are capable of offering the individual services required by the customers. The entire set of vehicles with different service qualifications available at the warehouse is denoted $M = \{1, 2, \ldots, m, \ldots, |M|\}$. Set M covers the disjointed subsets of vehicles M_k, $k \in K$, corresponding to the service types offered to the customers. It is assumed that a sufficiently large number of vehicles of each vehicle type ($|M_k| \geq |N|$) are available at the beginning of the planning horizon. Each vehicle type is associated with a type-specific cost rate c^k_{vehicle} due to the differences in both driver qualification and equipment required. Determining the number of vehicles used per type is part of the decision-making process. Moreover, customer orders have to be assigned to the vehicles used and sequenced. Variables u_{im}, $i \in N$, $m \in M$, determine whether a customer order i is assigned to vehicle m ($u_{im} = 1$) or not ($u_{im} = 0$). The sequencing decision, that is, the vehicle route, is quantified by decision variable x_{ij}^m, $i, j \in N^0$, $m \in M$. The variable quantifies whether customer location i is visited immediately before location j by vehicle m ($x_{ij}^m = 1$) or not ($x_{ij}^m = 0$). Each route starts and ends at the warehouse that is indexed with 0, $N^0 = N \cup \{0\}$. Vehicles have a uniform weight $\overline{w}$ and space $\overline{v}$ capacity that must not be violated when designing vehicle tours. Please note that specific packing restrictions are not considered.

In addition, a cost is incurred relating to the distances driven, that is, $c_{ij}^k, i, j \in N^0, k \in K$. Independent of its type, each vehicle can perform one tour at most. Before a tour can start, all items included in the customer orders assigned must be prepared by the warehouse, that is, loaded onto the vehicle. The earliest starting time of vehicle m is denoted by b_m. In addition, the driving times between locations are denoted by $t_{i,j}$, $i, j \in N^0$.

Table 1 summarizes the sets, parameters and decision variables defined. The **OAS-VRPSD** is modeled afterwards.

$$minimize \sum_{k \in K} \sum_{m \in M_k} \sum_{i \in N^0} \sum_{j \in N^0} c_{ij}^k \cdot x_{ij}^m + \sum_{k \in K} \sum_{m \in M_k} \sum_{i \in N} c_{\text{vehicle}}^k \cdot x_{0i}^m + \sum_{i \in N} c_{\text{picker}} \cdot y_i \quad (1)$$

subject to

$$y_i + \sum_{j \in N, i \neq j} z_{ji} = 1 \quad \forall i \in N \quad (2)$$

$$\sum_{j \in N^+, i \neq j} z_{ij} = 1 \quad \forall i \in N \quad (3)$$

$$r_i \geq p_i \quad \forall i \in N \quad (4)$$

$$r_i \geq r_j + p_i - Q(1 - z_{ji}) \quad \forall i, j \in N, i \neq j \quad (5)$$

$$b_m \geq r_i - Q \cdot (1 - u_{im}) \quad \forall i \in N, m \in M \quad (6)$$

$$\sum_{m \in M} u_{im} = 1 \quad \forall i \in N \quad (7)$$

$$u_{im} \leq e_{ik} \quad \forall i \in N, m \in M_k, k \in K \quad (8)$$

$$u_{jm} = \sum_{i \in N^0, i \neq j} x_{ij}^m \quad \forall j \in N, m \in M \quad (9)$$

$$u_{jm} = \sum_{i \in N^0, i \neq j} x_{ji}^m \quad \forall j \in N, m \in M \quad (10)$$

$$\sum_{i \in N} x_{0i}^m \leq 1 \quad \forall m \in M \quad (11)$$

$$a_i \geq b_m + t_{0i} - Q \cdot (1 - u_{im}) \quad \forall i \in N, m \in M \quad (12)$$

$$a_j \geq a_i + t_{ij} + s_i - Q \cdot (1 - x_{ij}^m) \quad \forall i, j \in N, i \neq j, m \in M \quad (13)$$

$$a_i \geq \alpha_i \quad \forall i \in N \quad (14)$$

$$a_i \leq \beta_i \quad \forall i \in N \quad (15)$$

$$\overline{v} \geq \sum_{i \in N} v_i \cdot u_{im} \quad \forall m \in M \quad (16)$$

$$\overline{w} \geq \sum_{i \in N} w_i \cdot u_{im} \quad \forall m \in M \quad (17)$$

$$x_{ij}^m \in \{0, 1\} \quad \forall i, j \in N^0, m \in M \quad (18)$$

$$u_{im} \in \{0, 1\} \quad \forall i \in N^0, m \in M \quad (19)$$

$$y_i \in \{0, 1\} \quad \forall i \in N \quad (20)$$

$$z_{ij} \in \{0, 1\} \quad \forall i \in N, j \in N^+, i \neq j \quad (21)$$

The objective function (1) quantifies the total costs that need to be minimized. The first term of the cost function specifies the traveling costs associated with all tours generated. The second term quantifies the vehicle employment costs, which depend on the number of vehicles used. Note that the cost parameters in both terms depend on the respective vehicle type deployed. Finally, the third term specifies the number of order pickers used and the corresponding picking costs. Constraints (2) to (5) represent the order picking subproblem. Constraints (6) connect order picking with vehicle routing, which is formulated in constraints (7) to (17). Definitions (18) to (21) specify the respective domains of the binary decision variables.

The release time of a delivery order, that is, the finishing date of a picking order, depends on the sequence of picking orders assigned to an order picker. An order i can either be processed first by one of the order pickers used ($y_i = 1$) or is the successor of another order (see Eq. (2)). Consequently, the order is either processed last ($z_{i,|N|+1} = 1$) or has at least one successor (see Eq. (3)). Inequations (4) and (5) define the finishing time of a picking order, which equals the release date of a delivery order (r_i).

The order picking and vehicle routing subproblems are connected by Inequalities (6). A tour cannot start until all the corresponding orders have been supplied by the warehouse. Within the vehicle routing subproblem, Eq. (7) ensure that a customer is assigned to exactly one vehicle. Inequations (8) ensure that the type of the vehicle fits the service qualification required. Note that the vehicle type chosen arises from the customer order(s) with the highest hierarchical service type of all orders assigned to the same tour. Equations (9) and (10) ensure the integrity of routing and customer-vehicle assignment variables (Irnich, Toth, & Vigo, 2014). They also guarantee that the depot is included in a route containing at least one customer. Constraints (11) specify that a vehicle performs one tour at most. Inequations (12) to (15) then define the earliest possible starting time of a service at customer's home, a_i, which cannot be earlier than when the vehicle arrives at that location or earlier than the lower bound of the delivery time window defined. In addition, the service is not allowed to take place after the upper bound of

TABLE 1 Notation for model OAS-VRPSD

Sets	
N	Set of customers, $N = \{1, 2, \ldots, i, \ldots, \lvert N \rvert\}$
N^0	Set of customers with 0 as warehouse, $N^0 = N \cup \{0\}$
N^+	Set of customers with $(\lvert N \rvert + 1)$ as an artificial customer, $N^+ = N \cup \{\lvert N \rvert + 1\}$
M	Set of vehicles, $M = \{1, 2, \ldots, m, \ldots, \lvert M \rvert\}$
K	Set of vehicle types, $K = \{1, 2, \ldots, k, \ldots, \lvert K \rvert\}$
M_k	Set of vehicles of type $k \in K$, $M_k \subseteq M$; $M_h \cap M_l = \emptyset\ \forall h, l \in K, h \neq l$
Parameters	
c_{ij}^k	Nonnegative travel cost between location i and j, if $i, j \in N^0$, is performed by vehicle type k, $k \in K$
c_{vehicle}^k	Fixed cost per vehicle used of type k, $k \in K$
c_{picker}	Fixed cost per picker used
e_{ik}	Binary parameter indicating whether customer i, $i \in N$, can be served by vehicle type k, $k \in K$, $(e_{ik} = 1)$ or not $(e_{ik} = 0)$
t_{ij}	Travel time between location i and j, $i, j \in N^0$
p_i	Pick time for the order of customer i, $i \in N$
s_i	Service time for the order of customer i, $i \in N$
v_i	Space requirement for the order of customer i, $i \in N$
$\overline{v}$	Maximum capacity of a vehicle, measured in space units
w_i	Weight of customer i's order, $i \in N$
$\overline{w}$	Maximum capacity of a vehicle, measured in weight units
Q	Sufficiently large number, for example, maximum length of planning horizon, that is, latest possible starting time of all services, $Q = \max_{i \in N}[\beta_i]$
α_i	Lower bound of the time window of customer i, $i \in N$
β_i	Upper bound of the time window of customer i, $i \in N$
Decision variables	
a_i	Starting time of service at customer i, $i \in N$
b_m	Starting time of vehicle m, $m \in M$, at the warehouse
r_i	Release date of the delivery order of customer i, $i \in N$
u_{im}	Binary variable that indicates whether customer i, $i \in N$, is assigned to vehicle m, $m \in M$, $(u_{im} = 1)$ or not $(u_{im} = 0)$
x_{ij}^m	Binary variable that indicates whether vehicle m, $m \in M$, visits location j, $j \in N^0$ immediately after location i, $i \in N^0$, $(x_{ij}^m = 1)$ or not $(x_{ij}^m = 0)$
y_i	Binary variable that indicates whether customer order i, $i \in N$, is processed first by an order picker $(y_i = 1)$ or not $(y_i = 0)$
z_{ij}	Binary variable that indicates whether the order of customer i, $i \in N$, is processed immediately before the order of customer j, $j \in N^+$, $i \neq j$, $(z_{ij} = 1)$ or not $(z_{ij} = 0)$

the defined delivery time window. The capacity limits of the vehicles also have to be considered (see constraints (16) and (17)).

The order picking subproblem is known to be NP-complete (Hochbaum & Shmoys, 1988) and the vehicle routing subproblem generalizes the NP-hard capacitated vehicle routing problem (Vidal et al., 2013a). Consequently, the more general **OAS-VRPSD** is NP-hard, too. Heuristic procedures are thus required to solve problem instances of practically relevant sizes.

5 | A GENERAL VARIABLE NEIGHBORHOOD SEARCH APPROACH

This section proposes a general variable neighborhood search algorithm that is adapted to the OAS-VRPSD. The VNS metaheuristic was introduced by Mladenović and Hansen (1997). A GVNS extends the basic metaheuristic by a more sophisticated local search procedure, that is, a variable neighborhood descent (VND) algorithm (Hansen, Mladenović, Brimber, & Pérez, 2010), while additionally considering the specific requirements of OAS-VRPSD problem characteristics. Although its concept is rather simple, the VNS metaheuristic has yielded several high-performing algorithms, for example, for vehicle routing problems with multiple periods (see Pirkwieser and Raidl (2008), Hemmelmayr, Doerner, and Hartl (2009) and Kovacs, Golden, Hartl, and Parragh (2014)) or multiple depots (see Salhi, Imran, and Wassan (2014)), and for order picking problems, for example, an order batching, batch assignment and sequencing problem (see Henn (2015)). In the subsequent section we first introduce a construction heuristic that generates a feasible solution for OAS-VRPSD. The GVNS approach developed uses this solution as a start for its improvement process. The construction heuristic solves the order picking and vehicle routing problem in a sequential manner, which is common in retail practice and in commercial parcel deliveries. Later we use this procedure as a benchmark when analyzing the GVNS developed within our numerical study.

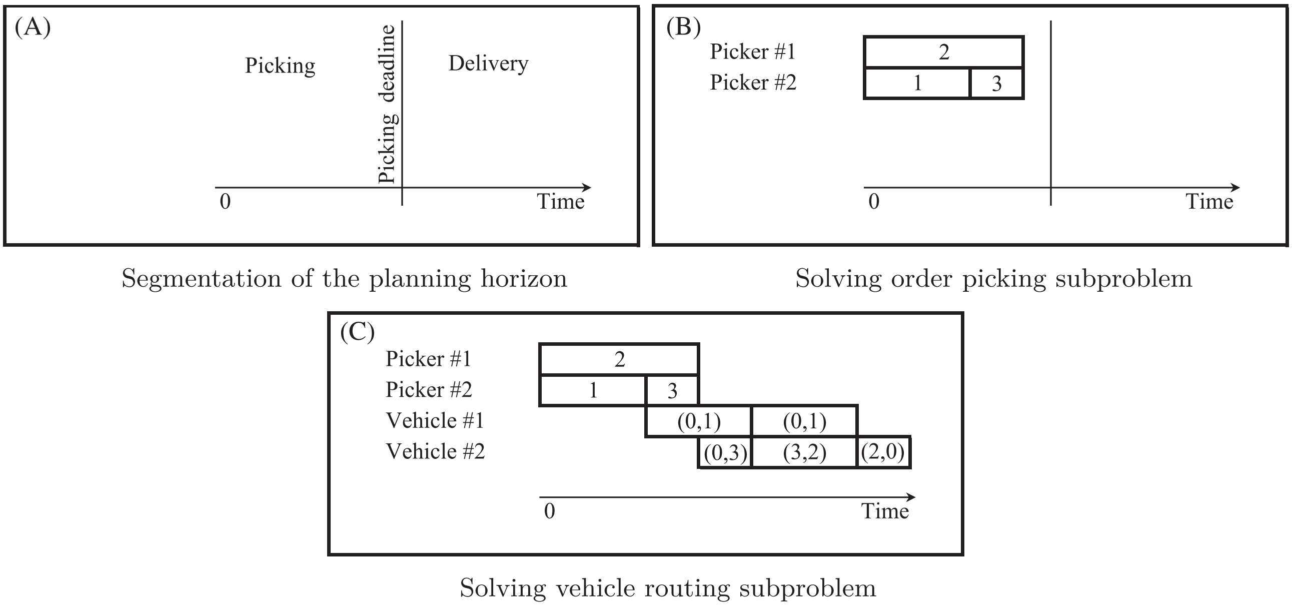

FIGURE 3 Scheme for the construction of an initial solution

5.1 | Initial solution

Figure 3 schematically depicts the general concept of the heuristic for building an initial solution for the problem being considered. The procedure splits the planning horizon (3a) into two planning phases according to partial workloads that have to be handled in each subproblem. We use the average upper time window as indicator for the length of the planning horizon. The sum of pick times and direct driving times defines the order picking and vehicle routing workload, respectively. The order picking subproblem is solved first by LPTR so that the release dates of the customer orders are determined (3b). Orders are assigned to the order pickers according to decreasing processing times, which balances the workload between the order pickers and minimizes their number. It has to be ensured for each assignment that the picking deadline is not violated and that the corresponding order can be delivered without violating the upper bound of its time window. If this is not possible, the order is assigned to an additional order picker, that is, the number of order pickers is increased by one. A solution to the vehicle routing subproblem is subsequently generated with the well-known savings algorithm of Clarke and Wright (1964). The release dates of the orders are taken as a fixed input (3c) in the savings algorithm. However, the original savings method does not consider different vehicle types, so they are set after the savings algorithm has terminated according to the highest service type of the customers assigned to a vehicle. This results in a feasible solution to the OAS-VRPSD. The integrated solution approach uses this solution as initiation of its improvement process.

5.2 | Integrated solution procedure

We apply a GVNS algorithm to obtain an integrated solution to the OAS-VRPSD problem. A GVNS consists of a shaking phase that perturbs the incumbent solution and a local search that derives benefits from the shaking moves, while neighborhood structures are changed in a systematic manner. The problem structure of the OAS-VRPSD leads, however, to special requirements for an integrated solution approach.

If one assumes a given solution to the vehicle routing problem, then the remaining cost only derives from the number of order pickers required in the order picking solution. The assignment of orders to pickers and the sequencing of picking orders has no impact on the target value. These factors do however determine the vehicle start dates via order release dates and thus whether a route remains feasible regarding the hard time window bounds. Following this argumentation, typical moves of local searches such as shift or swap either do not change the objective value—while maintaining feasibility—or lead to a non-feasible solution to the overall problem if they are applied to the order picking solution.

The GVNS applied considers these requirements of the OAS-VRPSD by only including vehicle routing decisions in the local search. Order picking decisions (together with resource and routing decisions) are then evaluated in the shaking phase. As a result, we directly evaluate all changes to the picking schedule to determine whether they have potential to improve the routing plan or not. In the course of doing this we immediately consider the consequences of a change on both subproblems. Algorithm 1 shows the main structure of our i_GVNS algorithm.

We assume a given sequence of neighborhood structures $\Phi_1^{GVNS}, \dots, \Phi_{max\_GVNS}^{GVNS}$ and an incumbent solution σ. The GVNS then randomly selects a neighbor solution $\widetilde{\sigma}$ from the ϕth neighborhood Φ_ϕ^{GVNS} of σ ($\Phi_\phi^{GVNS}(\sigma)$) in the shaking phase, and then applies a local search procedure to $\widetilde{\sigma}$. A new best solution always updates the incumbent solution. If no new best solution is found but $\widetilde{\sigma}$ complies with an acceptance criterion, $\widetilde{\sigma}$ replaces σ as the incumbent solution and the first neighborhood structure will be used in the shaking phase of the next iteration. Otherwise, σ remains unchanged and the next neighborhood structure is deployed. If the acceptance criterion is not met in an iteration of GVNS with the last neighborhood structure $\Phi^{GVNS}_{max\_GVNS}$ and no termination criterion has been met, the GVNS again applies the first neighborhood structure for the next iteration. Instead of a simple local search, we embed a VND as the local search procedure in our algorithm (Hansen et al., 2010).

Algorithm 1 i_GVNS

Input: problem data, number of neighborhood structures *max_GVNS* and *max_VND*
Output: best solution σ_{best} and corresponding total costs $f(\sigma_{best})$ determine (initial) solution σ

$$\sigma := VND(\text{problem data}, \overline{\Phi}^{VND}, \sigma)$$

while termination criterion not met **do**
 $\phi := 1$
 while $\phi \leq max\_GVNS$ **do**
  generate $\widetilde{\sigma}$ randomly from the ϕth neighborhood of σ ($\widetilde{\sigma} \in \Phi^{GVNS}_{\phi}(\sigma)$) // shaking
  $\widetilde{\sigma} := VND(\text{problem data}, \overline{\Phi}^{VND}, \widetilde{\sigma})$
  if $f(\widetilde{\sigma}) < f(\sigma_{best})$ **then**
   $\sigma_{best} := \widetilde{\sigma}$
   $\sigma := \widetilde{\sigma}$
   $\phi := 1$
  else if acceptance criterion met **then**
   $\sigma := \widetilde{\sigma}$
   $\phi := 1$
  else
   $\phi := \phi + 1$
  end if
  if σ_{best} not improved for certain number of iterations **then**
   $\sigma := \sigma_{best}$
  end if
 end while
end while

5.2.1 | VND procedure

The local search procedure of the GVNS, that is, the VND algorithm, follows the idea of the (G)VNS metaheuristic but changes neighborhood structures in a deterministic manner. In contrast to a (G)VNS, the entire neighborhood is explored and the best neighbor solution $\widetilde{\sigma}$ is determined. If the objective function value of $\widetilde{\sigma}$ ($f(\widetilde{\sigma})$) constitutes an improvement regarding the objective function value of σ, $\widetilde{\sigma}$ is set as the new incumbent solution and VND applies the first neighborhood structure again. Otherwise the next neighborhood of σ will be examined in the next iteration. The VND terminates, if the last neighborhood structure $\Phi^{VND}_{max\_VND}$ has been investigated but no improvement of σ results. Consequently, σ provides a local optimum to all neighborhoods of σ ($\Phi^{VND}_1(\sigma), \ldots, \Phi^{VND}_{max\_VND}(\sigma)$).

5.2.2 | Neighborhood structures in the local search

The VND approach contains two types of neighborhood structure Φ^{VND}_{ϕ}, that is, the Or-Opt Operator (Or, 1976) and the Relocate Operator (Savelsbergh, 1992).

- *Or-Opt Operator*: As an intra-route move, the Or-Opt Operator shifts one customer or a sequence of customers to another position in the same route.
- *Relocate Operator*: As an inter-route move, the Relocate Operator shifts one customer or a sequence of customers to a different route that can be an existing route or a newly created one. While the Or-Opt Operator has no impact on the vehicle type, the Relocate Operator also balances between routing cost and costs for vehicle types.

5.2.3 | Neighborhood structures in the shaking phase

In the shaking phase of the GVNS algorithm, neighborhood structures are associated with the vehicle routing subproblem, order picking subproblem, and the determination of resources, that is, increasing or reducing the number of pickers or vehicles.

- *Cross-exchange operator*: Two sequences of customers of different routes are interchanged (Taillard, Badeau, Gendreau, Guertin, & Potvin, 1997). The lengths of both sequences may be different and are determined independently within identical limits. This operator also makes it explicitly possible to swap customers to a new vehicle and to adjust vehicle types according to the customers exchanged. Similar to the Relocate Operator, the Cross-exchange Operator takes into account both routing costs and costs per vehicle type.
- *Shift same picker operator*: A customer order or a sequence of orders is moved to a different position at the same picker.
- *Shift different pickers operator*: A customer order or a sequence of orders is moved to a different picker.
- *Swap different pickers operator*: Customer orders or sequences from two different pickers are exchanged. Similar to the Cross-exchange Operator, the lengths of both sequences considered are determined independently but follow the same limit.
- *New picker operator*: A number of randomly chosen orders are selected from set $\{\underline{\omega}, \ldots, \overline{\omega}\}$ and assigned to a newly added picker. The orders can originate from different pickers and are assigned to the newly added picker using the sequence selected.

- *Delete picker operator*: A randomly chosen picker is deleted and the corresponding orders are randomly assigned to random positions of the remaining pickers.
- *Delete vehicle operator*: A randomly chosen vehicle (tour) is deleted and the corresponding orders are randomly assigned to the remaining delivery tours.

Please note that vehicle capacities are always respected but a shaking move to the order picking problem may result in a non-feasible solution regarding time window compliance. The local search then aims to restore feasibility by only considering feasible moves. The acceptance criterion rejects the current solution if no feasible solution is found.

5.2.4 | Acceptance and termination criteria

Particularly for vehicle routing problems, threshold acceptance (Dueck & Scheuer, 1990) has proven to be an effective mechanism to avoid neighborhood search-based algorithms from getting stuck in local optima (see, e.g., Tarantilis, Kiranoudis, and Vassiliadis (2004), Polacek, Hartl, and Doerner (2004), or Henke, Speranza, and Wäscher (2015)). Basically, an altered solution $\widetilde{\sigma}$ is accepted if the corresponding objective value $f(\widetilde{\sigma})$ is smaller than the objective value $f(\sigma)$ of the incumbent solution σ plus a threshold T. First, a new best solution is always accepted and the incumbent solution is updated as well. Second, we implement a threshold-acceptance procedure depending on the objective value of the current best solution σ_{best}. In doing this, the threshold T is determined by $\gamma \cdot f(\sigma_{best})$. A solution $\widetilde{\sigma}$ is then accepted as a (new) incumbent solution σ if $f(\widetilde{\sigma}) < f(\sigma) + T$. This procedure achieves slightly better results than only considering the objective value of the incumbent solution. The parameter γ is dynamically adjusted within the algorithm. Initially, γ is set to zero such that only improvements are accepted within the search behavior. However, after a predefined number of iterations without finding a new incumbent solution, γ is increased by parameter μ with the result that solutions of inferior quality can be accepted. If a new incumbent or best solution has been found, γ is reset back to its initial value. The GVNS algorithm terminates if a certain time limit is reached.

6 | NUMERICAL EXPERIMENTS

This section presents the design and results of the numerical experiments that are performed to evaluate our solution approach for the OAS-VRPSD. In the following we will refer to our GVNS approach described in the previous section as integrated GVNS (i_GVNS). In the first subsection we outline the main benchmark approaches for evaluating the performance of i_GVNS (Section 6.1). Afterwards we detail the configuration of our test instances (Section 6.2), the implementation setting (Section 6.3) and describe the results based on instances generated in the basic setting (Section 6.4), as well as with varied planning scenarios (Section 6.5). We then briefly analyze the solution quality over time (Section 6.6) before finally testing i_GVNS in a real-life setting using a case from a major European electronics retailer (Section 6.7).

6.1 | Benchmark approaches

As a main benchmark approach (SEQ), a sophisticated algorithm is applied that solves the OAS-VRPSD sequentially as it is prevalent in retail practice. The approach selected corresponds as far as possible to the planning activities of our practice partner. It can therefore be assumed that the SEQ approach will lead to comparable results as a direct application of approaches currently used in retail practice. The SEQ approach proceeds as follows. The initial solution is determined by the construction heuristic described in Section 5.1. Orders are therefore assigned to a number of order pickers that has previously been estimated using the LPTR. Solving the order picking subproblem applying LPTR minimizes the number of order pickers and thus the decision-relevant costs in the order picking subproblem. Subsequently, an adapted savings approach is used to solve the vehicle routing problem taking into account all constraints. To obtain a fair comparison for our integrated solution approach and to minimize the decision-relevant costs arising in the vehicle routing subproblem, the initial solution is improved, first via the local search procedure of the i_GVNS algorithm, that is, a VND approach. Second, a GVNS algorithm is applied to the vehicle routing subproblem. This means that all neighborhood structures that deal with the order picking subproblem are removed ($\Phi_4^{GVNS} - \Phi_{13}^{GVNS}$ and Φ_{G15}^{VNS}) from i_GVNS, maintaining the sequential planning characteristics in retail practice. SEQ therefore uses intelligence from our i_GVNS while keeping separate both planning domains being considered. A comparison of the solution quality of the i_GVNS approach (integrated solution) with the benchmark (sequential solution) thus shows whether—and if so, to what extent—benefits arise from dealing with the OAS-VRPSD as a simultaneous planning problem. Both i_GVNS and SEQ are performed with the same computing time and parameter settings in order to have a fair comparison (see Section 6.3). Moreover, an exact approach is applied by implementing the model formulation from Section 4 in a standard optimization software. However, this approach is only suitable for solving small-sized instances.

6.2 | Test instances

We set up our experiments on newly generated problem instances in order to evaluate the corresponding contributing factors since the OAS-VRPSD has not been dealt with so far. Instances of previous studies would only have been reusable to a limited extent as they do not differentiate in terms of pick times and size of the delivery area, nor do they build on soft upper time windows, order batching or multiple vehicle tours.

However, the specifications of the characteristics are strongly oriented to practical data.

Our research focus lies on same-day deliveries for large products, so the planning horizon is set equal to a few hours only. More precisely, if a hypothetical upper bound for the latest possible delivery date for all customers is set (eg, 9 PM), the length of the planning horizon determines the point in time by which customers can order such that they are served the same day. In our experiments, the length of the planning horizon is set to 6, 9, and 12 hours, that is, customers can order by 3 PM, noon, and 9 AM, respectively. The planning horizon is thus described as tight (t), medium (m) or wide (w).

We generate problem instances with 5, 7, 50, 100, and 200 customers. Instances with 5 and 7 customers are used to compare the performance of the i_GVNS approach to the exact one, while instances of practically relevant sizes are the basis for the evaluation of the i_GVNS compared to the sequential planning approach, SEQ. The number of orders for practically relevant instances is oriented to what we find at the partner company as well as on problem sizes typically used in literature. Henn (2015) and Scholz et al. (2017), for example, use problem instances with up to 200 customer orders for order picking problems, Desaulniers, Madsen, and Ropke (2014) state that about 150 customer orders are realistic problem sizes for vehicle routing problems, and Schubert et al. (2018) generated instances with up to 200 customer orders for an integrated order picking and vehicle routing problem adaptable to the daily deliveries of supermarkets from a central DC.

The pick times of customer orders depend on several factors, although a discrete picking strategy is assumed, for example, the size of the warehouse, the storage assignment policy, the picker routing strategy, etc. (De Koster et al., 2007). We generate small (s), medium (m) and large (l) pick times randomly within certain limits, that is, {[5, 10], [10, 30], [30, 50]} minutes, to be able to evaluate the influence of pick times in our setting.

The customers and the DC are randomly located within a predefined delivery area. The size of the delivery area is adapted from Holzapfel, Hübner, Kuhn, and Sternbeck (2016), Ullrich (2013), and Schubert et al. (2018) and set at 50×50, 100×100, and 200×200 km. These areas represent “Metropolitan” to “District” regions (Holzapfel et al., 2016). Driving times and distances between the corresponding locations are calculated according to the Euclidean distances. Correspondingly, the size of the delivery area (SDA) is characterized as small (s), medium (m), or large (l).

A customer order includes a time window within which the corresponding service must start. Each customer is assigned to a specific time window with a length of 2 hours. The positions of these time windows are randomly chosen within the planning horizon. More precisely, the lower bound of a time window is at least as great as the sum of the pick time of the corresponding order and the driving time from the DC to the respective customer location. The problem generator ensures that a problem instance is feasibly solvable and that the length of a time window is not reduced by the end of the planning horizon. That means the upper bound of a time window is strictly lower than or equal to the hypothetical upper bound for the latest possible delivery.

TABLE 2 Domains of problem class parameters

Problem class parameter		Domain
Number of customer orders	$\|N\|$	5, 7, 50, 100, 200
Pick times	PT	[5, 10], [10, 30], [30, 50] minutes
Tightness of planning horizon	TPH	6, 9, 12 hours
Size of delivery area	SDA	50×50, 100×100, 200×200 km

We assume that a customer is equally likely to choose one of three service types. The timespans correspond to those applicable in practice. For service type 1 (all installations and connection services), a service time of 80 minutes is defined. Service type 2 (simple installations) is processed in 40 minutes, while service type 3 (delivery to place of destination) requires 20 minutes. Three vehicle types are considered according to these service types. A vehicle of the first type can be deployed for 840 monetary units (MU) and used to serve all service types. Vehicles of the second type cost 680 MU and can be used for service types two and three, while a vehicle of the last type (type 3) can only serve customers who order service three, and costs 580 MU. The cost factor per kilometer driven is set to 1 MU. Finally, an order picker can be deployed in the warehouse for 240 MU. The cost parameters are generally inspired by the values from the case study and from previous collaborations with industry partners. The weight and space requirement of the items ordered by a customer are randomly generated within [20, 120] kg and [60×60, 120×120] cm^2. Each vehicle can load up to 1,500 kg, and its loading area has a size of 14.64 m^2.

In total, 135 problem classes result from the combination of problem characteristics mentioned above. Ten instances have been generated for each problem class. Table 2 shows the corresponding problem class parameters and their domains.

6.3 | Implementation and parameter testing

The experiments were carried out on Intel Xenon E5-2697 v3 processors. IBM ILOG CPLEX Optimization Studio version 12.6.3 (CPLEX) was used for the implementation of the model formulation from Section 4. The i_GVNS algorithm and the sequential benchmark approach (SEQ) are programmed in C++ and compiled with the g++ -O3 optimization.

Table 3 shows the specification of the neighborhood structures applied in the local search (Φ_ϕ^{VND}) and their sequence of application in the VND. In total, VND makes use of six neighborhood structures.

Fifiteen neighborhood structures, Φ_ϕ^{GVNS}, are embedded in the shaking phase of the GVNS algorithm. Table 4 provides an overview of the operators considered and parameters

TABLE 3 Embedded neighborhood structures in the local search

ϕ	Operator	Sequence length
1	Or-Opt	1
2	Or-Opt	2
3	Or-Opt	3
4	Relocate	1
5	Relocate	2
6	Relocate	3

TABLE 4 Embedded neighborhood structures in the shaking phase

		Minimum	Maximum
ϕ	Operator	Sequence length	
1	Cross-exchange	1	1
2	Cross-exchange	1	2
3	Cross-exchange	1	3
4	Shift same picker	1	1
5	Shift same picker	1	2
6	Shift same picker	1	3
7	Shift different pickers	1	1
8	Shift different pickers	1	2
9	Shift different pickers	1	3
10	Swap different pickers	1	1
11	Swap different pickers	1	3
12	Swap different pickers	1	5
		Number of orders	
13	New picker	2	10
14	Delete vehicle		
15	Delete picker		

used. The number of orders assigned to a picker who is additionally deployed, that is, $\underline{\omega}$ and $\overline{\omega}$, is set to 2 and 10, respectively.

i_GVNS and SEQ terminate if a certain time limit has been reached. For instances with 5 and 7 orders, the algorithm was executed for 1 second. In order to solve instances with 50, 100, and 200 orders, the CPU limit was set to 5, 10, and 30 minutes, respectively. In contrast, CPLEX was not bounded by a CPU limit to solve the instances considered to optimality.

We tested different parameter constellations regarding the sequence length and number of orders used within the operator tests. Accordingly, we tested the parameters used within the threshold accepting procedure. In total we pretested 50 instances with random problem parameters chosen based on the problem classes described in Section 6.2. A further increase in the neighborhood structures did not lead to significant improvements. Moreover, we found that modifying threshold T after 60 iterations without a new incumbent solution results in an appropriate tradeoff between intensification and diversification. In detail, we modify T by increasing γ by $\mu = 0.01$, that is, T is increased by 1% of the best solution. Moreover, the incumbent solution was reset to the best solution after 600 iterations without finding a new best solution.

6.4 | Comparison of solution approaches

In the following we present the results from our test instances based on data generated. First, the results of the small instances are discussed comparing i_GVNS to an exact solution of the planning problem (Section 6.4.1). Second, i_GVNS and SEQ are compared to each other using the remaining instances with a large number of customer orders (Section 6.4.2).

6.4.1 | Results of the experiments with small instances

This subsection compares the performance of i_GVNS and SEQ with the performance of an exact approach, that is, implementation of the model formulation from Section 4 in CPLEX. Table 5 presents the corresponding results (columns 4 to 13). The first three columns specify the problem class, that is, the pick times (PT), the size of the delivery area (SDA), and the tightness of the planning horizon (TPH) in accordance with Section 6.2. The average deviation of the target values of the solutions generated by the i_GVNS from the optimal target value (dev) is stated in columns 4 and 9 for five and seven customers, respectively. Columns 5 and 10 show the improvement of SEQ by i_GVNS, that is, imp_seq $= (f(\sigma^{\text{SEQ}}) - f(\sigma^{\text{i\_GVNS}}))/f(\sigma^{\text{SEQ}})$, with $f(\sigma^{\text{i\_GVNS}})$ and $f(\sigma^{\text{SEQ}})$ as final solutions of i_GVNS and SEQ, respectively. Moreover, the computing time in seconds after which the best solution is found is given in columns 6 and 11 for the exact approach (t_ex), and in columns 7 and 12 for the i_GVNS algorithm (t_GVNS). Finally, the number of optimally solved instances per problem class by the i_GVNS (#opt) is shown in columns 8 and 13.

i_GVNS is able to optimally solve 530 out of 540 instances, that is, 98.15%, with a maximum deviation from the optimal objective value of 1.09%. The solution quality of the i_GVNS is only 0.0048% worse than the optimal target value, on average. For five customer orders, the exact approach is able to generate optimal solutions for all instances within 0.25 seconds on average, while the i_GVNS requires less than 0.01 seconds until the final solution has been found. However, the average time to find an optimal solution with the exact approach increases to 24.63 seconds in the case of seven customers. The i_GVNS algorithm still finds a final solution within fractions of a second (0.03 seconds on average). i_GVNS performs approximately 25,600 and 24,000 iterations within the computing time of 1 second for instances with five and seven customer orders, respectively. Moreover, the computational times of the exact approach strongly vary even within the same problem class. The maximum computing time of the exact approach amounts to 1957.73 seconds (large pick times, small delivery square, tight planning horizon). In

TABLE 5 Evaluation of the solution quality of i_GVNS for small and large problem instances compared to benchmark approaches

			5 orders					7 orders					50 orders	100 orders	200 orders
PT	SDA	TPH	dev	imp_seq	t_ex	t_GVNS	#opt	dev	imp_seq	t_ex	t_GVNS	#opt	imp_seq	imp_seq	imp_seq
s	s	t	0.03	0.00	0.27	0.01	8	0.00	1.30	5.26	0.07	10	4.56	3.68	2.21
s	s	m	0.00	0.00	0.21	0.00	10	0.00	0.00	62.34	0.01	9	4.06	2.74	2.46
s	s	l	0.00	0.00	0.20	0.00	10	0.00	0.00	4.49	0.00	10	4.31	5.31	3.57
s	m	t	0.00	1.36	0.22	0.00	10	0.00	7.58	3.57	0.00	9	6.89	5.29	3.65
s	m	m	0.00	0.00	0.20	0.00	10	0.01	1.26	1.47	0.01	9	5.96	3.62	3.20
s	m	l	0.00	2.23	0.22	0.00	10	0.00	1.25	0.69	0.00	10	2.87	3.76	3.28
s	l	t	0.00	8.75	0.24	0.00	10	0.00	12.61	5.84	0.00	10	13.56	12.65	13.05
s	l	m	0.00	8.27	0.22	0.00	10	0.00	9.59	16.98	0.00	10	11.97	12.98	13.48
s	l	l	0.00	6.13	0.22	0.00	10	0.00	8.80	4.39	0.00	10	10.80	10.67	12.65
m	s	t	0.01	3.56	0.22	0.00	9	0.00	6.42	10.76	0.02	10	13.69	13.17	10.04
m	s	m	0.00	0.08	0.22	0.00	10	0.02	0.75	2.51	0.04	9	12.75	12.60	7.57
m	s	l	0.00	0.00	0.23	0.00	10	0.01	2.00	2.57	0.00	9	10.97	11.00	5.98
m	m	t	0.00	3.08	0.29	0.01	10	0.00	4.77	8.04	0.04	10	12.16	8.02	6.56
m	m	m	0.00	0.02	0.30	0.00	10	0.00	4.46	3.72	0.01	10	9.09	7.62	4.63
m	m	l	0.00	0.75	0.29	0.03	10	0.10	1.80	3.22	0.03	9	8.80	7.16	4.59
m	l	t	0.00	8.68	0.33	0.00	10	0.07	9.61	28.28	0.03	9	7.80	6.63	4.61
m	l	m	0.00	9.66	0.24	0.01	10	0.00	9.47	6.41	0.03	10	6.47	5.22	3.60
m	l	l	0.00	6.06	0.25	0.00	10	0.00	8.00	20.66	0.03	10	6.54	4.27	2.91
l	s	t	0.00	11.49	0.34	0.02	10	0.00	19.67	308.04	0.09	10	17.10	13.26	12.24
l	s	m	0.00	7.66	0.31	0.01	10	0.00	21.94	10.15	0.08	10	14.63	11.32	8.71
l	s	l	0.00	4.54	0.31	0.00	10	0.00	15.33	7.53	0.02	10	17.73	10.82	7.19
l	m	t	0.00	11.81	0.36	0.03	10	0.00	20.49	47.68	0.08	10	19.78	15.36	12.45
l	m	m	0.00	5.08	0.36	0.01	10	0.00	16.08	9.82	0.01	10	14.96	12.20	7.82
l	m	l	0.00	0.00	0.20	0.00	10	0.00	23.63	9.11	0.04	10	14.28	12.68	7.16
l	l	t	0.00	8.43	0.27	0.01	10	0.00	10.27	24.38	0.09	10	11.05	10.27	6.69
l	l	m	0.00	7.29	0.23	0.00	10	0.00	4.70	32.33	0.09	10	8.99	7.65	5.31
l	l	l	0.00	6.46	0.24	0.00	10	0.00	12.92	24.89	0.03	10	8.56	8.45	4.00

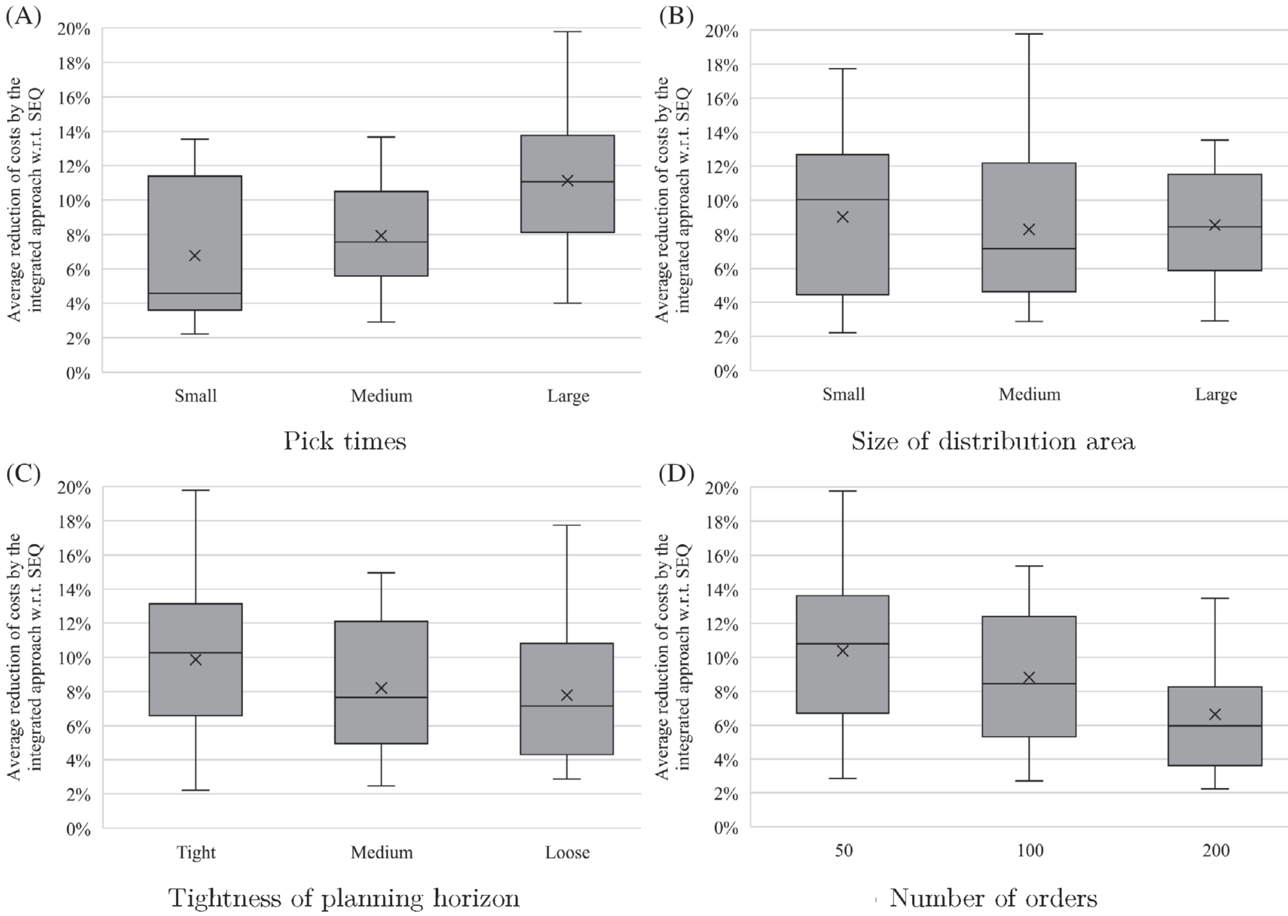

FIGURE 4 Impact of parameter values (attributes) for four problem parameters on the cost benefits between i_GVNS and SEQ

the same problem class, the minimum time to find an optimal solution amounts to 0.95 seconds. Comparing the objective values of SEQ and i_GVNS reveals 6.59% lower target values for i_GVNS solutions on average considering the cases of five and seven orders. The differences in solution quality between SEQ and i_GVNS are greater for seven (8.69%) than for five (4.5%) orders, since the coordination effort necessary increases from five to seven orders. Moreover, the solution quality of SEQ strongly varies between different problem classes. For example, both approaches generate a similar solution quality for five orders and small pick times and a small distribution area, but a deviation of about 19% can be observed for seven orders, large pick times and a small distribution area. Similar results occur for large instances. We therefore refer to the following section that analyzes the results of SEQ and i_GVNS in more detail.

6.4.2 | Results of the experiments with large instances

The following part of the analysis investigates the benefits of an integrated solution approach for the OAS-VRPSD. The solution quality of the benchmark approach SEQ is compared to i_GVNS for instances of practically relevant sizes for this purpose.

Table 5 contains the results for large instances with 50, 100, and 200 customer orders (columns 14–16) solved with SEQ and i_GVNS, respectively. The table shows the average deviation of the objective function values of SEQ compared to the corresponding target values of i_GVNS (imp_seq).

i_GVNS outperforms SEQ by 8.62% across all problem instances. The maximum average percentage of improvement per problem class equals 19.77% (50 orders, PT = large, SDA = medium, TPH = tight), while the minimum average percentage of improvement is equal to 2.21% (200 orders, PT = small, SDA = small, TPH = tight). The solutions of the integrated approach require an average of 2.0 pickers and 1.8 vehicles less than the solutions of the sequential approach. Using a lower number of vehicles increases the average number of stops per vehicle. In detail, vehicles of type 1 deliver 0.5 customer orders more in i_GVNS solutions than in SEQ solutions, that is, 5.63 compared to 5.13, while vehicles of type 2 deliver 0.29 customer orders more in that constellation, that is, 6.62 compared to 6.33. Both approaches, however, lead to similar results for type 3 vehicles, that is, the average number of stops equals 3.78 and 3.74, respectively. Summarizing these results, the SEQ approach assigns slightly fewer customer orders to the vehicles of each type than the i_GVNS approach. However, the distribution of stops at a customer type by a vehicle type is very similar for both approaches. Overall, the routing costs are reduced by 8.02%.

We analyzed the percentage of improvement of i_GVNS compared to SEQ in respect of three values (attributes) for four problem parameters, that is, pick times, size of distribution area, tightness of planning horizon, and number of orders. Figure 4 shows the results by means of box plots for each of these parameters and respective values (attributes). The center line and the cross depict the median and mean, respectively.

Increasing pick times favors an integrated planning approach (see Figure 4A). The percentage of improvement of i_GVNS to SEQ increases from 6.79% to 7.94%, and 11.14% on average for small, medium, and large pick times, respectively. The average number of orders that can be handled by an order picker within a certain amount of time decreases with increasing pick times, and thus the average release date of the delivery orders increases. Consequently, if the order pickers release delivery orders later, coordination between order picking and vehicle routing becomes more important since there is less time available to arrange the delivery tours in an appropriate manner.

Varying the size of the distribution area (SDA) does not seem to affect the percentage of improvement of i_GVNS to SEQ (see Figure 4B). The results amount to 9.02%, 8.29%, and 8.55% for small, medium, and large distribution areas, respectively. However, the impact of the size of the delivery area impacts the percentage of improvement of i_GVNS to SEQ in combination with the pick times. SEQ approximates the number of order pickers by the ratio of the sum of pick times and the direct driving times from the DC to the customers. Using i_GVNS solutions as benchmark, this calculation scheme works quite well for similar combinations of pick times and driving times, that is, small/small, medium/medium, and large/large. Both approaches apply a similar number of order pickers in these scenarios, that is, SEQ applies 3% fewer pickers. However, if travel times are considerably more pronounced than pick times, SEQ overestimates the necessary number of pickers by 38%, and underestimates them by 22% if travel times are considerably less pronounced than pick times. An overestimated number of pickers leads to routes that are 0.76% shorter in SEQ solutions, while an underestimated number results in routes that are 27.89% longer. For similar resources, vehicle routes determined by SEQ are 10.98% longer compared to routes calculated by i_GVNS. This result shows the importance of coordinated scheduling of picking and delivery resources in the environment considered, which is achieved by i_GVNS.

Moreover, the tightness of the planning horizon impacts the percentage of improvement between i_GVNS and SEQ (see Figure 4C). In the case of a loose planning horizon the limited capacities pose a greater challenge than meeting the delivery time windows. However, if the planning horizon is tightened, the composition of tours becomes more restricted by the release dates and the time windows of the customer orders, and an integrated solution to the OAS-VRPSD becomes more important as a result. In our experiments, we only consider very tight planning horizons with a length of 6, 9, and 12 hours and, additionally, hard time windows so that the time dimension is generally more restrictive than the capacity constraints. As a result, the impact of the tightness of the planning horizon shows a clear trend, although the differences between the average percentages of improvement of the respective attributes are relatively small. They amount to 9.87%, 8.21%, and 7.79% for a tight, medium or loose planning horizon. The results of the median values support this observation.

The average percentage of improvement of i_GVNS to SEQ decreases with the number of customers (see Figure 4D). However, the absolute savings increase by applying i_GVNS instead of SEQ. For 50, 100, and 200 orders, they amount to 1,338.25, 2,073.46, 2,947.25 MU on average per instance. In combination with the limited planning horizon, more pickers are necessary to handle the increased workload. This increases the number of orders picked per time unit that are ready for delivery. Consequently some delivery tours can start earlier since the number of customer orders per vehicle remains similar. Note that the number of customers per vehicle is physically limited by the capacity and especially due to the time window constraints so that the average number of orders per tour will only go up slightly if the entire number of customer orders increases. The sequential planning approach (SEQ) benefits relatively more from this situation than the integrated approach (i_GVNS). Nevertheless, the absolute advantages of the integrated approach are still present (see above). Ramaekers et al. (2018) observe a similar development for increasing instance sizes in an integrated picking and vehicle routing problem with time windows.

6.5 | Varied planning scenarios

As already mentioned, the most common delivery mode offered in online retailing is next-day delivery. In order to analyze the impact of different planning scenarios for same-day delivery we additionally extend the problem classes introduced in Section 6.2 according to the delivery conditions and number of vehicle types prepared. The results can be used to analyze the impact of different same-day delivery modes on logistics costs and to answer the question as to the conditions under which an integrated solution approach is especially worth applying.

6.5.1 | Setting and assumptions

Customers are accustomed to deliveries within customer-specific time windows. However, this significantly restricts planning and the solution space. We therefore introduced two additional delivery conditions, that is, two general time windows and a general delivery deadline by which the service of a customer must start. The planning horizon is split in order to assign customers to one of the two general time windows. Customers are randomly assigned to one of the two time windows on condition that the amount of time for serving a customer must not be shorter than 2 hours (length of a customer-specific time window). In the case of delivery deadlines we align the associated values to the length of the planning horizon, that is, 6, 9, and 12 hours, respectively.

Second, the vehicle fleet is modified on the assumption that the requirements of customers cannot be influenced with respect to the service type demanded. It must therefore be ensured that the compilation of the vehicle fleet is compiled such that all service types can be performed. However, the deployment of more than one vehicle type increases the complexity of the planning problem at the operational level as well as in contract design with subcontractors or for the compilation of a proper vehicle fleet, respectively. For this reason, only the vehicles with the highest service level (type 1) are prepared in a first variation, while in a second the first and second highest qualified vehicle types (types 1 and 2) are prepared.

In order to obtain valid results, we use the identical instances as for the previous experiments, but modify the corresponding delivery conditions for the customers and/or the number of vehicle types. All problem instances are then solved with i_GVNS under the same conditions, and the corresponding solutions are compared to analyze the cost impact of the varied scenarios.

6.5.2 | Impact of delivery conditions on costs

The total decision-relevant costs decrease by 18.20% if customer-specific time windows are excluded such that customers are to be served by a general delivery deadline. These cost savings result from a reduction of resources and from significantly fewer driving distances. First, the routing costs decrease by 22.38%, which follows from the fact that the compilation of tours is no longer constrained by time windows. Second, the number of order pickers is reduced by 3.36 on average, resulting in 19.65% lower (decision-relevant) order picking costs. Third, the number of vehicles required decreases on average by 3.68 that leads on average to 17.40% lower decision-relevant costs. The reduction in the deployment of resources is explained by the significant time savings due to the tour compilation.

If two general time windows are established, the results are similar to those cited above. The total costs decrease by 10.23%, the vehicle-dependent costs by 11.45%, and the routing costs by 13.31% for all problem instances. However, the number of order pickers deployed increases by 0.89 (5.53%) on average, as half of the customers have to be served very early within the planning horizon.

6.5.3 | Impact of number of vehicle types on costs

Table 6 summarizes the average number of resources that are deployed within the corresponding instances of problem classes with one, two, and three vehicle types, analyzing the benefits of a hierarchically divided vehicle fleet. The total costs increase if the number of vehicle types available are reduced. The percentage of additional costs amounts to 0.68% and 4.74%, respectively, if only the two highest or the highest classes of vehicles can be deployed. Since the lowest class of vehicles is only rarely used (1.1 vehicles on average), the amount of additional costs if this type is abandoned would be fairly low. This is not true for vehicles of type 2 that are used at a rate of 29.21% and 36.05% in problem classes with three and two vehicle types, respectively. From the results, a tradeoff for the deployment of different resources can be identified. If only the highest class of vehicles is deployed, route planning is easier and more cost efficient and the number of order pickers can be reduced. On the other hand, vehicle costs increase by 1.57% and 8.53% if only the two highest or the highest vehicle type can be deployed.

6.5.4 | Changes regarding the benefits of an integrative solution to the OAS-VRPSD

As mentioned in Section 6.4.2, the time window constraints greatly restrict the solution space due to the short planning horizon. The number of customers per tour is fairly small as a result. For the extended experiments above, the benefits of an integrated solution to the OAS-VRPSD rise significantly since the number of customers per tour increases, and it is therefore harder for the DC to ensure tours. If different delivery scenarios (general time windows, general deadline) are considered as well, the percentage of the improvement using an integrated solution approach increases from 8.62% to 12.98% on average. The percentage of improvement amounts to 12.12% for problem classes with two general delivery time windows. For a general delivery deadline, the benefits increase to the remarkable figure of 18.19%.

In contrast to the delivery conditions, the number of vehicle types prepared does not have a significant impact on the benefits of an integrated solution to the OAS-VRPSD.

6.6 | Solution quality over time

For next-day deliveries, computing times of one or several hours are not a critical issue since customer orders are known since the previous evening at the latest (Schubert et al., 2018). However, this is not true for same-day deliveries due to the tightness of the planning horizon in this setting. This section evaluates i_GVNS regarding the computational times.

Figure 5 shows the average deviation from the best target value found (ADBTV) for all problem classes with 50, 100, and 200 customer orders. These instances were solved for 5, 10, and 30 minutes, respectively, in order to deal with the increasing complexity arising from a larger number of orders. The initial solution of i_GVNS features a target value that is still far above the best solution found (31.14%, 29.91%, 27.85%). The target value however significantly decreases within the first seconds of runtime. In detail, ADBTV amounts to 4.83%, 7.32%, and 8.80% for 50, 100, and 200 orders, respectively, after 1% of the computational times. After 20% of the computational times, ADBTV varies

TABLE 6 Deployment of resources and compilation of costs

			No. vehicles			Costs		
No. vehicle types	f(i_GVNS)	No. pickers	Type 1	Type 2	Type 3	Pickers	Vehicles	Routing
Three	21,018.3	16.5	12.1	5.4	1.1	3,955.9	14,458.5	2,604.0
Two	21,161.4	16.3	12.0	6.8	0.0	3,903.9	14,670.6	2,586.9
One	22,013.8	16.0	18.7	0.0	0.0	3,847.5	15,692.4	2,473.9

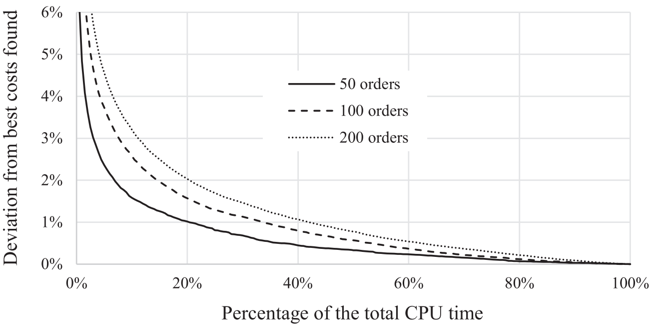

FIGURE 5 Development of solution quality as a function of applied computational time

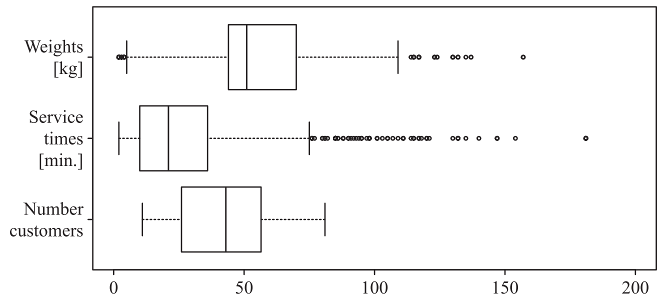

FIGURE 6 Boxplots of the number of customers, service times (minutes) and weights (kg) in the case study. Please note that four outliers for service times, that is, 206, 208, 257, and 279, are not depicted here

between 1.03% and 2.07% only. Moreover, i_GVNS still finds improvements within the last 20% of computational time, when the corresponding ADBTV accounts for 0.07%, 0.12% and 0.21%, respectively.

We perform additional experiments to analyze the long-term runtime behavior of i_GVNS. In doing so, we set the termination criterion of the algorithm to 6 hours when solving the problem instances with 100 and 200 customer orders. The additional computational time improves the previous best solution found by 2.19% and 2.15%, respectively. Also, ADBTV amounts to 0.32% and 0.52% after half of the computing time and, even after 5.4 hours, the i_GVNS still manages to slightly improve the incumbent solution by 0.03% (100 orders) and 0.06% (200 orders).

Summarizing these insights, the computational times applied are adequately chosen for all problem classes, but i_GVNS still manages to find improvements after 6 hours of computational time for large instances. However, in real-life settings there is a tradeoff between the improved solution quality of the application of the algorithm and the time for achieving this, which could be used for additional orders to be prepared for delivery. Overall, if computational times are very critical as in same-day delivery settings, i_GVNS is appropriate as it manages to improve the objective value significantly within the first seconds of application.

6.7 | Real-life case

In this final subsection of our analysis we present the results of the application of i_GVNS and SEQ to real-life data. The data is provided by an internationally operating European electronics retailer and covers historical customer data of approximately 2 months including more than 2,500 customer orders from one regional DC. This corresponds to approximately 50 customer orders per day, which have to be fulfilled within 6 hours. Out of this data, 47 instances are generated that meet the basic expectations of the retailer regarding same-day customer orders for the corresponding area. The resulting instance size is a multiple of the size that can reasonably be solved by implementation of the model formulation from Section 4. Consequently, both i_GVNS and SEQ are applied to each instance for 15 minutes. A direct comparison of the method used in practice was unfortunately not possible for technical and confidential reasons. However, in terms of conception and algorithmics, the SEQ approach corresponds as far as possible to the current planning process at our case company.

The distribution of the number of customer orders per day and the corresponding service times (minutes) and weights (kg) are shown in a boxplot in Figure 6.

Note that volume data has not been provided, and is therefore not considered in the case study. However, consultations show that volume is a critical issue that is currently monitored by dispatchers. Pick times are uniformly distributed in a range of [5, 15] minutes. The size of the distribution area based on the customer locations equals 146×81 km. Finally, the ratio between the average vehicle cost rate and the cost rate for order pickers equals 2.5.

The results of the application of i_GVNS and SEQ are provided in Table 7 in the form of normalized values.

The objective values of the solutions generated by SEQ are 18.9% above those of i_GVNS, which demonstrates the necessity of an integrated solution approach to the OAS-VRPSD in practice. The corresponding cost savings of the i_GVNS solutions result from reducing the deployment of all types of resources. While the reduction in the number

TABLE 7 Percentage savings when applying i_GVNS compared to SEQ in the real-life case

Δtotal cost	ΔNo. pickers	ΔNo. vehicles				Δrouting costs
		Total	Type 1	Type 2	Type 3	
18.9%	3.8%	23.2%	13.5%	37.5%	14.7%	17.2%

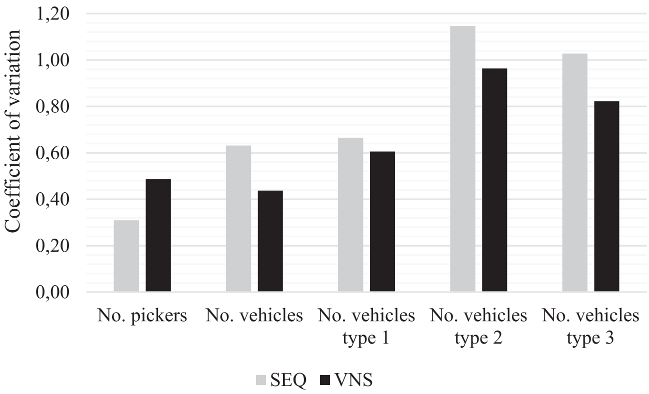

FIGURE 7 Coefficients of variation for the deployment of resources in the real-life case; SEQ: gray bars, i_GVNS: black bars

of pickers appears to be relatively small, at 3.8% on average, the deployment of vehicles is reduced by a remarkable 23.2%. Moreover, savings related to vehicle routing amount to 17.2% using the i_GVNS approach.

Besides cost savings potential, we investigate the variations of the resources deployed by both the sequential and integrated solution approach. The provision of a sufficiently large vehicle fleet and picking team is a natural component of medium- and long-term contracts with subcontractors and employees. Variations in the resources required must also be taken into account in this context. A larger variation in order pickers can be more easily compensated for as their numbers can also be increased by using temporary employment agencies or similar. A greater variation in vehicles, on the other hand, is more difficult to handle, since a truck must be provided in addition to (trained) drivers. The size and contractual fixed costs of the vehicle fleet are thus determined depending on a fixed service level regarding vehicle provision. The coefficients of variation (CV) of the deployed resource types are depicted in Figure 7 for all 47 instances.

Generally, the CVs are relatively high, which can be explained by a high CV for the number of customers (0.44). However, there are significant differences for individual resource types and solution approaches. SEQ (gray bar) results in a lower CV for order pickers (0.31) compared to the total number of vehicles (0.63). This is not true for i_GVNS (black bar), which leads to a lower CV for the total number of vehicles (0.44) compared to the CV of the number of order pickers (0.49). This is interesting from a practical point of view. One aspect of note is that vehicle and vehicle routing costs are the largest cost drivers. Another is that a two-man team, the vehicle, and the corresponding equipment has to be allocated to prepare a vehicle. It therefore appears easier to handle a higher variation in the necessary order pickers required than vehicles such that the results—besides the cost savings potential—indicate a more robust solution structure for our integrated solution approach.

Figure 8 details the dependency of resources applied depending on the number of orders considered. Circles (vehicles) and squares (pickers) indicate the number of resources applied. The lines show the third degree polynomial trend line for the number of pickers required and the exponential trend line for the number of vehicles needed. The number of pickers and vehicles used increases synchronously with the number of orders when applying the integrated planning approach (i_GVNS). In sequential planning (SEQ), however, the number of vehicles required grows strongly as the number of orders rises, while the number of pickers increases only slightly. This is due to the fact that in the sequential approach picking is optimized first and vehicle routing plans afterwards. Thus, the VRP solution has to build on the picking schedule already fixed. The results show that this leads to a fundamental disadvantage of the sequential planning approach compared to i_GVNS. The resources planned increase almost linearly with the number of orders considered in i_GVNS solutions, while the number of vehicles planned increases exponentially in the SEQ solution in the case study. This is also an indicator of a more robust solution of i_GVNS and its benefits for application in practice.

The CVs of the different vehicle types indicate that the volatility of resources applied is higher for the cheaper and less flexible vehicle types (types 2 and 3). This means that in practice the retailer can make plans based on more stable resource requirements using higher-cost vehicle types, but needs flexibility especially regarding lower-cost vehicle types. Note that i_GVNS generates solutions with a lower CV compared to SEQ for all vehicle types (see Figure 7).

The results of the case study confirm the results obtained using the problem instances generated. An integrated solution approach results in significant cost savings for both practical and simulated data. Moreover, the case study shows that solutions of the integrated approach are easier to handle in the long run due to a smaller variation in the number of vehicles, but at the expense of a higher variation in the number of order pickers.

In the following we vary the cost rate per order picker (CRPP, i.e., parameter c_{picker}) to obtain additional insights on the solution structure. In so doing, the 47 practical problem instances described above (basic instances) are solved again, but the CRPP is set to 25%, 50%, 200% and 400% of the current rate. Table 8 shows the corresponding results of i_GVNS normalized to the solutions of the basic instances.

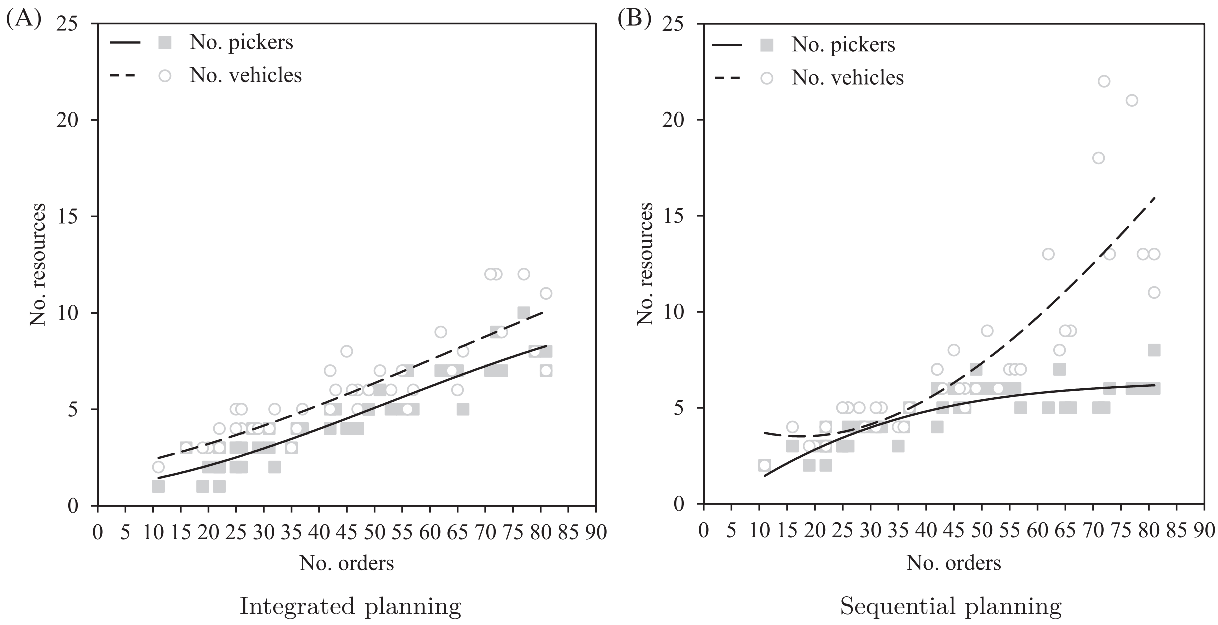

FIGURE 8 Resources applied depending on the number of orders considered

TABLE 8 Results for varied cost rates per order picker planned

Multiplicator picking costs	0.25	0.50	1.00	2.00	4.00
Target value	0.83	0.89	1.00	1.21	1.51
Number of pickers	1.12	1.07	1.00	0.99	0.80
Number of vehicles	0.98	0.98	1.00	1.02	1.12
Ratio vehicles to pickers	1.14	1.22	1.32	1.35	1.93

The reduction in order picking costs, that is, CRPP, naturally leads to a reduction in total costs and vice versa. Decreasing CRPP to 25% leads to solutions with 16.72% lower target values, while a reduction of 50% results in 11.21% lower target values. However, the target value increases greatly, that is, by 21.39% and 51.26%, if the CRPP is doubled and quadrupled, respectively. These cost increases are induced by the higher cost rates per picker but also by changes to the solution structure.

Reducing the CRPP to 25% increases the number of order pickers planned by 12.18%, and reducing it to 50% leads to a 6.87% increase. The savings in target value are obtained by reducing the number of vehicles used and creating improved routes. More order pickers available reduce the average order release date by 7.52% for a CRPP equal to 25% of the initial CRPP, and 5.09% for a CRPP of 50%. Consequently, the number of vehicles used goes down by 2.33% and 2.11% and routing costs by 2.34% and 0.65%, respectively. The same argument holds for an increased CRPP, leading to fewer order pickers planned, more vehicles used, and ultimately increasing target values. The increase in terms of number of vehicles, however, results mainly from an increase in vehicles of type 3. It increases by 41.10% comparing the smallest and largest CRPP. In contrast, the number of vehicles of type 2 increases by only 11.21%, while the number of vehicles of type 1 goes up by the much lower figure of 2.29%.

Summarizing this analysis, the ratio of the number of vehicles to number of pickers used is heavily affected by CRPP (see Table 8). This ratio is rather balanced for the smallest CRPP but close to two for a quadrupled CRPP, which shows the tradeoff between number of pickers and number of vehicles.

7 | CONCLUSIONS AND OUTLOOK ON FUTURE RESEARCH

7.1 | Conclusions

This paper considers the simultaneous solution of order picking and order delivery where the delivery of products includes the possibility of product-associated installation services at the customer's home. The decision problem at hand is especially relevant when distributing large durable consumer goods in a same-day delivery environment, and is observable in the case of both pure online and omnichannel retailers. The problem integrates the order assignment and sequencing problem with a vehicle routing problem with site dependencies. We denote the entire problem OAS-VRPSD. The paper presents a decision support model and an integrative solution approach that is based on a generalized variable neighborhood search (GVNS) algorithm. The integrative solution approach considerably improves the solution of a sequential approach that is the current standard in retail practice. The paper undertakes an extensive numerical study for both simulated data and real-life data from a leading European omnichannel retailer, which impressively confirms the value of integration and reveals several managerial insights:

(a) The logistics costs are highly dependent on the delivery conditions offered. Offering two general time windows or a general delivery deadline instead of accepting customer-specific time windows reduces costs by 10.23% and 18.20%, respectively, in our numerical study. This effect becomes especially visible when solving both problems in an integrative manner. More delivery flexibility for the retailer therefore requires an integrative solution approach.

(b) The logistics costs also depend on the number of different vehicle types used. The cost differences become most visible when operating only one instead of two vehicle types. However they become negligible when operating only two instead of three vehicle types. Operating two instead of three vehicle types therefore seems appropriate in the environment considered by the generated data. However, no special effect has been observed on the necessity of an integrative solution approach. The integrative approach shows equal benefits compared to the sequential approach in all of these parameter variations.

(c) The integrative approach leads to more robust (stable) results related to the number of vehicles used than the sequential approach. This is especially true for the vehicle type offering the most value-added services. The integrative solution approach therefore offers the opportunity for the retailer in our case study to operate with a relatively constant mix of vehicle types even though the customer orders significantly vary from day to day.

(d) The analysis of the cost rate per order picker (CRPP) in the case study leads to additional insights into the development of the solution structure. The analysis reveals the tradeoff between the number of pickers and the number of vehicles. In addition, the same analysis shows that an increase in picking cost is mostly compensated by the least flexible, that is, the lowest-cost vehicle type. This is a valuable hint for any retailer since this type of vehicle is much easier to enlarge or reduce than more flexible types if demand and/or cost rates change.

7.2 | Outlook on future research

The suggested modeling and solution approach includes several specific—possibly restricting—assumptions, leading to new challenges for further research:

(a) The timeline of the modeling approach assumes that all processes—including the picking process—are only allowed to start after the cut-off time. It may be of interest to know under what conditions this assumption could be relaxed, and how this would influence the results achieved.

(b) This research assumes a deterministic planning environment. This enables us to identify unbiased insights and interconnections in the problem being considered. Furthermore, the approach developed can be used as a planning tool that achieves better results than classical sequential approaches. However, the very tight timeline can intensify the effects caused by uncertainties in the picking process and during delivery. Take for example the cases where a delivery arrives at the customer location only after the time window scheduled due to an unforeseen extension of travel times or an extended installation process at a previous customer. The current approach can only take this into account by providing sufficient slack in the pick times, travel times, service times and time windows. However, slack times that are too generous result in unnecessary costs. Investigation of the effects of uncertainties in an integrated problem setting, especially in relation to the length of the planning horizon, would provide a valuable opportunity for further research.

(c) We assume a sufficient number of available order pickers and vehicles per type in order to adapt the resources required in the short term. This assumption is appropriate for large retailers, as in our case, and comes at high expense for this flexibility by contract. However, if the number of resources is restricted, a feasible solution that serves all customers is no longer ensured.

(d) We assume a sufficient number of loading docks at the DC. Again, this is true for large DCs, on which our research is based. For smaller DCs, vehicle loading operations have to be scheduled, so additional decisions have to be made, and further interactions are involved.

(e) Omnichannel grocery retailers serve supermarkets and online customers from joint regional DCs. Retailers operate multi-compartment vehicles in order to deal with diverse product segments that require individual temperature requirements, (Hübner & Ostermeier, 2019). A branch of further research could therefore concentrate on integrating the order picking process of diverse segments and the joint supply of these segments using multi-compartment vehicles.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the computer and data resources provided by the Leibniz Supercomputing Centre (www.lrz.de), and the doctoral scholarship from the Erich-Kellerhals-Stiftung (Foundation), Ingolstadt, Germany, for one of the authors. In addition, we would like to thank the anonymous reviewers and the Guest Editors of the special issue for their valuable recommendations, which have significantly improved our paper. Open access funding enabled and organized by Projekt DEAL.

ORCID

Heinrich Kuhn https://orcid.org/0000-0003-3704-4042

REFERENCES

Agatz, N. A. H., Fleischmann, M., & van Nunen, J. A. E. E. (2008). E-fulfillment and multi-channel distribution—A review. *European Journal of Operational Research*, *187*(2), 339–356.

Archetti, C., Feillet, D., & Speranza, M. G. (2015). Complexity of routing problems with release dates. *European Journal of Operational Research*, *247*, 797–803.

Archetti, C., Jabali, O., & Speranza, M. G. (2015). Multi-period vehicle routing with due dates. *Computers & Operations Research*, *61*, 122–134.

bevh. (2019). *Warengruppen im online-Handel nach dem Umsatz in Deutschland in den Jahren 2015 bis 2018*. Retrieved from https://de.statista.com/statistik/daten/studie/253188/

Biskup, D., Herrmann, J., & Gupta, J. N. D. (2008). Scheduling identical parallel machines to minimize total tardiness. *International Journal of Production Economics*, *115*, 134–142.

Cattaruzza, D., Absi, N., & Feillet, D. (2016). The multi-trip vehicle routing problem with time windows and release dates. *Transportation Science*, *50*, 676–693.

Chen, Z. (2010). Integrated production and outbound distribution scheduling: Review and extension. *Operations Research*, *58*, 130–148.

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

Cordeau, J.-F., Gendreau, M., & Laporte, G. (1997). A tabu search heuristic for periodic and multi-depot vehicle routing problems. *Networks*, *30*, 105–119.

Cordeau, J.-F., & Laporte, G. (2001). A tabu search algorithm for the site dependent vehicle routing problem with time windows. *INFOR*, *39*, 292–298.

Cordeau, J.-F., Laporte, G., & Mercier, A. (2004). Improved tabu search algorithm for the handling of route duration constraints in vehicle routing problems with time windows. *Journal of the Operational Research Society*, *55*, 542–546.

Cordeau, J.-F., & Maischberger, M. (2012). A parallel iterated tabu search heuristic for vehicle routing problems. *Computers & Operations Research*, *39*, 2033–2050.

De Koster, R., Le-Duc, T., & Roodbergen, K. J. (2007). Design and control of warehouse order picking: A literature review. *European Journal of Operational Research*, *182*, 481–501.

Delorme, M., Iori, M., & Martello, S. (2016). Bin packing and cutting stock problems: Mathematical models and exact algorithms. *European Journal of Operational Research*, *255*(1), 1–20.

Desaulniers, G., Madsen, O. B. G., & Ropke, S. (2014). *The vehicle routing problem with time windows*. In P. Toth & D. Vigo (Eds.), Vehicle routing: Problems, methods and applications (2nd ed., pp. 119–160). Philadelphia, PA: Society for Industrial and Applied Mathematics and the Mathematical Optimization Society.

Dueck, G., & Scheuer, T. (1990). Threshold accepting: A general purpose optimization algorithm appearing superior to simulated annealing. *Journal of Computational Physics*, *1*, 161–175.

Gallino, S., & Moreno, A. (2019). Operations in an Omnichannel world—Springer series in supply chain management (Vol. *8*). Berlin: Springer International Publishing.

Hansen, P., Mladenović, N., Brimber, J., & Pérez, J. A. M. (2010). *Variable neighborhood search*. In M. Gendreau & J.-Y. Potvin (Eds.), Handbook of metaheuristics International Series in Operations Research & Management Science (Vol. *146*, 2nd ed., pp. 61–86). New York, NY et al.: Springer.

HDE. (2019). *Umsatz durch E-Commerce (B2C) in Deutschland in den Jahren 1999 bis 2018 sowie eine Prognose für 2019*. Retrieved from https://de.statista.com/statistik/daten/studie/3979/

Hemmelmayr, V. C., Doerner, K. F., & Hartl, R. F. (2009). A variable neighborhood search heuristic for periodic routing problems. *European Journal of Operational Research*, *195*, 791–802.

Henke, T., Speranza, M. G., & Wäscher, G. (2015). The multi-compartment vehicle routing problem with flexible compartment sizes. *European Journal of Operational Research*, *246*, 730–743.

Henn, S. (2015). Order batching and sequencing for the minimization of the total tardiness in picker-to-part warehouses. *Flexible Services and Manufacturing*, *27*, 86–114.

Hochbaum, D. S., & Shmoys, D. B. (1988). A polynomial approximation scheme for machine scheduling on uniform processors: Using the dual approximating approach. *SIAM Journal on Computing*, *17*, 539–551.

Holzapfel, A., Hübner, A., Kuhn, H., & Sternbeck, M. G. (2016). Delivery pattern and transportation planning in grocery retailing. *European Journal of Operational Research*, *252*, 54–68.

Holzapfel, A., Kuhn, H., & Sternbeck, M. G. (2018). Product allocation to different types of distribution center in retail logistics networks. *European Journal of Operational Research*, *264*, 948–966.

Hübner, A., Holzapfel, A., & Kuhn, H. (2015). Operations management in multi-channel retailing: An exploratory study. *Operations Management Research*, *8*(3), 84–100.

Hübner, A., Holzapfel, A., & Kuhn, H. (2016). Distribution systems in omni-channel retailing. *Business Research*, *9*(2), 255–296.

Hübner, A., Holzapfel, A., Kuhn, H., & Obermair, E. (2019). Distribution in Omni-channel grocery retailing: An analysis of concepts realized. Series in supply chain management (pp. 283–309). Berlin: Springer.

Hübner, A., & Ostermeier, M. (2019). A multi-compartment vehicle routing problem with loading and unloading costs. *Transportation Science*, *53*(1), 282–300. https://doi.org/10.1016/j.cor.2016.12.023

Irnich, S., Toth, P., & Vigo, D. (2014). *The family of vehicle routing problems*. In P. Toth & D. Vigo (Eds.), Vehicle routing—Problems, methods and applications, MOS-SIAM Series on Optimization; No. 18 (2nd ed. SIAM, Philadelpia, pp. 1–33).

Klapp, M. A., Erera, A. L., & Toriello, A. (2018). The dynamic dispatch waves problem for same-day delivery. *European Journal of Operational Research*, *271*, 519–534.

Kovacs, A. A., Golden, B. L., Hartl, R. F., & Parragh, S. N. (2014). The generalized consistent vehicle routing problem. *Transportation Science*, *49*, 796–816.

Kuhn, H., Schubert, D., & Holzapfel, A. (2020). Integrated order batching and vehicle routing – a general adaptive large neighborhood search algorithm. *European Journal of Operational Research*, available online 10 April 2020.

McKinsey. (2017). *Warenzustellung am selben Tag vor dem Durchbruch*. Retrieved from https://www.mckinsey.de/warenzustellung-am-selben-tag-vor-dem-durchbruch

Mladenović, N., & Hansen, P. (1997). Variable neighborhood search. *Computers & Operations Research*, *24*, 1097–1100.

Moons, S., Braekers, K., Ramaekers, K., Caris, A., & Arda, Y. (2019). The value of integrating order picking and vehicle routing decisions in a B2C e-commerce environment. *International Journal of Production Research*, *57*(20), 6405–6423.

Moons, S., Ramaekers, K., Caris, A., & Arda, Y. (2017). Integrating production scheduling and vehicle routing decisions at the operational decision level: A review and discussion. *Computers & Industrial Engineering*, *104*(2), 224–245.

Moons, S., Ramaekers, K., Caris, A., & Arda, Y. (2018). Integration of order picking and vehicle routing in a B2C context. *Flexible Services and Manufacturing Journal*, *30*, 813–843.

Morihara, I., Ibaraki, T., & Hasegawa, T. (1982). Bin packing and multiprocessor scheduling problems with side constraint on job types. *Discrete Applied Mathematics*, *6*, 173–191.

Or, I. (1976). *Traveling salesman-type combinatorial problems and their relation to the logistics of regional blood banking*. (Ph.D. thesis). Northwestern University, Evanston, IL.

Paul, J., Agatz, N., Spliet, R., & Koster, R. D. (2019). Shared capacity routing problem—An omni-channel retail study. *European Journal of Operational Research*, *273*(2), 731–739.

Pirkwieser, S., Raidl, G. R. (2008). A variable neighborhood search for the periodic vehicle routing problem with time windows. In: F. Troyes (Ed.), *Proceedings of the 9th EU Meeting on Metaheuristics for Logistics and Vehicle Routing*.

Pisinger, D., & Ropke, S. (2007). A general heuristic for vehicle routing problems. *Computers & Operations Research*, *34*, 2403–2435.

Polacek, M., Hartl, R. F., & Doerner, K. (2004). A variable neighborhood search for the multi depot vehicle routing problem with time windows. *Journal of Heuristics*, *10*, 613–627.

Ramaekers, K., Caris, A., Moons, S., & van Gils, T. (2018). Using an integrated order picking-vehicle routing problem to study the impact of delivery time windows in e-commerce. *European Transport Research Review*, *10*(56), 1–11.

Salhi, S., Imran, A., & Wassan, N. A. (2014). The multi-depot vehicle routing problem with heterogeneous vehicle fleet: Formulation and a variable neighborhood search implementation. *Computers & Operations Research*, *52*, 315–325.

Savelsbergh, M. W. P. (1992). The vehicle routing problem with time windows: Minimizing route duration. *ORSA Journal on Computing*, *4*, 146–154.

Schmid, V., Doerner, K. F., & Laporte, G. (2013). Rich routing problems arising in supply chain management. *European Journal of Operational Research*, *224*, 435–448.

Scholz, A., Schubert, D., & Wäscher, G. (2017). Order picking with multiple pickers and due dates—simultaneous solution of order batching, batch assignment and sequencing, and picker routing problems. *European Journal of Operational Research*, *263*, 461–478.

Schubert, D., Scholz, A., & Wäscher, G. (2018). Integrated order picking and vehicle routing with due dates. *OR Spectrum*, *40*(4), 1109–1139.

Taillard, E. D., Badeau, P., Gendreau, M., Guertin, F., & Potvin, J.-Y. (1997). A tabu search heuristic for the vehicle routing problem with soft time windows. *Transportation Science*, *31*, 170–186.

Tarantilis, C. D., Kiranoudis, C. T., & Vassiliadis, V. S. (2004). A threshold accepting metaheuristic for the heterogeneous fixed fleet vehicle routing problem. *European Journal of Operational Research*, *152*, 148–158.

Tompkins, J. A., White, J. A., Bozer, Y. A., & Tanchoco, J. M. A. (2010). Facilities planning (4th ed.). Hoboken, NJ: John Wiley & Sons.

Toth, P., & Vigo, D. (Eds.). (2014). Vehicle routing: Problems, methods and applications (2nd ed.). Philadelphia, PA: Society for Industrial and Applied Mathematics and the Mathematical Optimization Society.

Ullrich, C. A. (2013). Integrated machine scheduling and vehicle routing with time windows. *European Journal of Operational Research*, *227*, 152–165.

van Gils, T., Caris, A., Ramaekers, K., & Braekers, K. (2019). Formulating and solving the integrated batching, routing, and picker scheduling problem in a real-life spare parts warehouse. *European Journal of Operational Research*, *277*(3), 814–830.

van Gils, T., Ramaekers, K., Caris, A., & Cools, M. (2017). The use of time series forecasting in zone order picking systems to predict order pickers' workload. *International Journal of Production Research*, *55*, 6380–6393.

van Gils, T., Ramaekers, K., Caris, A., & de Koster, R. B. (2018). Designing efficient order picking systems by combining planning problems: State-of-the-art classification and review. *European Journal of Operational Research*, *267*(1), 1–15.

Vidal, T., Crainic, T. G., Gendreau, M., & Prins, C. (2013a). Heuristics for multi-attribute vehicle routing problems: A survey and synthesis. *European Journal of Operational Research*, *231*, 1–21.

Vidal, T., Crainic, T. G., Gendreau, M., & Prins, C. (2013b). A hybrid genetic algorithm with adaptive diversity management for a large class of vehicle routing problems with time-windows. *Computers & Operations Research*, *40*, 475–489.

Voccia, S. A., Campbell, A. M., & Thomas, B. W. (2017). The same-day delivery problem for online purchases. *Transportation Science*, *53*(1), 167–184.

Wollenburg, J., Holzapfel, A., Hübner, A., & Kuhn, H. (2018). Configuring retail fulfillment processes for omni-channel customer steering. *International Journal of Electronic Commerce*, *22*(4), 540–575.

Wollenburg, J., Hübner, A., Kuhn, H., & Trautrims, A. (2018). From bricks-and-mortar to bricks-and-clicks—Logistics networks in omni-channel grocery retailing. *International Journal of Physical Distribution & Logistics Management*, *48*(4), 415–438.

Zare-Reisabadi, E., & Mirmohammadi, S. H. (2015). Site dependent vehicle routing problem with soft time window: Modeling and solution approach. *Computers & Industrial Engineering*, *90*, 177–185.

How to cite this article: Schubert D, Kuhn H, Holzapfel A. Same-day deliveries in omnichannel retail: Integrated order picking and vehicle routing with vehicle-site dependencies. *Naval Research Logistics* 2021;68:721–744. https://doi.org/10.1002/nav.21954