European Journal of Operational Research xxx (xxxx) xxx

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/eor

Production, Manufacturing, Transportation and Logistics

Exact algorithms for routing electric autonomous mobile robots in intralogistics

Anne Meyer [a,*], Timo Gschwind [b], Boris Amberg [c], Dominik Colling [a]

[a] Karlsruhe Institute of Technology (KIT), Germany
[b] RPTU Kaiserslautern-Landau, Germany
[c] FZI Forschungszentrum Informatik, Germany

ARTICLE INFO

Keywords:
Routing
Autonomous mobile robots
Electric vehicle routing problem
Branch-price-and-cut
Recharging strategies

ABSTRACT

In intralogistics and manufacturing, autonomous mobile robots (AMRs) are usually electrically powered and recharged by battery swapping or induction. We investigate AMR route planning in these settings by studying different variants of the electric vehicle routing problem with due dates (EVRPD). We consider three common recharging strategies: battery swapping, inductive recharging with full recharges, and inductive recharging with partial recharges. Moreover, we consider two different objective functions: the standard objective of minimizing the total distance traveled and the minimization of the total completion times of transport jobs. The latter is of particular interest in intralogistics, where time aspects are of crucial importance and the earliest possible completion of jobs often has priority. In this context, recharging decisions also play an essential role. For solving the EVRPD variants, we propose exact branch-price-and-cut algorithms that rely on ad-hoc labeling algorithms tailored to the respective variants. We perform an extensive computational study to generate managerial insights on the AMR route planning problem and to assess the performance of our solution approach. The experiments are based on newly introduced instances featuring typical characteristics of AMR applications in intralogistics and manufacturing and on standard benchmark instances from the literature. The detailed analysis of our results reveals that inductive recharging with partial recharges is competitive with battery swapping, while using a full-recharges strategy has considerable drawbacks in an AMR setup.

1. Introduction

Advanced hardware and control software currently allows the introduction of autonomous mobile robots (AMRs) for intralogistics tasks, e.g., in manufacturing sites, warehouses, or transshipment terminals (Fragapane, de Koster, Sgarbossa, & Strandhagen, 2021). Conventional autonomous vehicles in intralogistics systems follow fixed tracks using lines, magnets, or barcodes on the floor or reflectors on the wall to determine their position (Furmans, Seibold, & Trenkle, 2019). The new generation of vehicles (typically referred to as AMRs) is free of infrastructure for localization and navigation (Furmans et al., 2019): AMRs move freely observing their environment with laser scanners or 3D cameras which enables them to avoid even dynamic obstacles. The low infrastructure requirements and high flexibility make AMRs successful on an industrial scale and, in many applications, outweigh the higher price per vehicle, which is mainly due to the expensive sensors.

AMRs are usually electrically powered and two relevant technologies for recharging exist: *battery swapping (BS)* and *inductive recharging (IR)*. In the former, used batteries are swapped with fully recharged ones. In the latter, the AMRs drive to a station equipped with a recharging mat, where they are recharged by induction. Battery swapping is usually faster but requires people or infrastructure for swapping. Inductive recharging is slower but cheaper as it does not require additional resources. Effective recharging decisions, i.e., *when*, *where* (at which recharging station), and *for how long* to recharge, are essential when solving electric vehicle routing problems (EVRP) (Schneider, Stenger, & Goeke, 2014). This is especially true in route planning for fleets of electric AMRs where these decisions are crucial for the performance of the system (Kabir & Suzuki, 2019; McHaney, 1995; Zou, Xu, De Koster, et al., 2018), because the battery capacity of AMRs is typically limited and several recharging stops are required throughout the day (Jun, Lee, & Yih, 2021; Le-Anh & De Koster, 2006). If BS is employed, the recharging decisions reduce to the questions of *when* and *where* to recharge as the time for swapping batteries is fixed assuming that the vehicle must not wait. The same is true for IR if a so-called *full-recharges strategy*

* Corresponding author.
E-mail addresses: anne.meyer@kit.edu (A. Meyer), gschwind@rptu.de (T. Gschwind), amberg@fzi.de (B. Amberg), dominik.colling@partner.kit.edu (D. Colling).

https://doi.org/10.1016/j.ejor.2024.12.041
Received 6 July 2023; Accepted 23 December 2024
Available online 31 December 2024

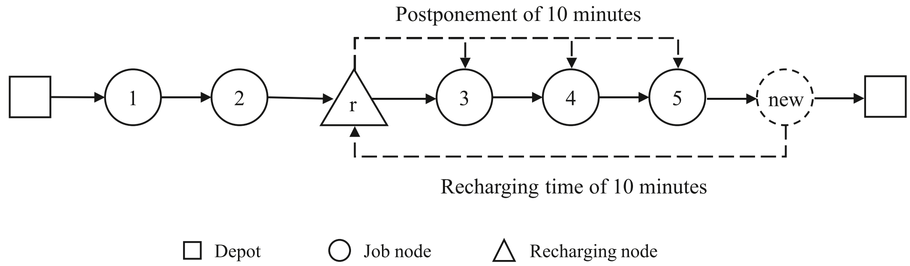

Fig. 1. Adding job node "new" to the tour results in a postponement of the completion times of stops 3, 4, and 5 in case of PR.

(FR) is applied, in which vehicles' batteries are always fully recharged whenever they visit a recharging station (Desaulniers, Errico, Irnich, & Schneider, 2016). In this case, the amount of energy to recharge and, thus, the recharging duration is directly determined by the battery level upon arrival at the station. In a so-called *partial-recharges strategy (PR)*, on the other hand, the vehicles are free to recharge any amount of energy when visiting a station so that the recharging durations are part of the decision (Desaulniers et al., 2016; Keskin & Çatay, 2016). Because it strongly simplifies the decision process, FR is often applied in the battery management of AMR fleets rather than the more flexible PR, although it can be detrimental for the performance of the system (De Ryck, Versteyhe and Shariatmadar, 2020).

In many intralogistics applications timing issues are crucial. Jobs usually have due dates that must be met and the primary objective is often to finish jobs as early as possible, i.e., to minimize completion times of jobs. The latter reduces the risk of disruptions caused by delays and makes vehicles available for future jobs and recharging. While the minimization of the sum of completion times is a common objective function of production scheduling problems (Pinedo, 2009, p. 130), it has received little attention in the literature on vehicle routing problems (VRPs) in general and in the EVRP literature in particular. The only exception we are aware of is the work of Jun et al. (2021) who propose a heuristic algorithm for a routing problem considering recharging decisions for AMR systems with a similar objective function, namely the minimization of the sum of tardinesses, i.e., the sum of the delays of all transport requests, also sometimes called the total tardiness. To the best of our knowledge, no exact method has been proposed considering recharging decisions, due dates, and the minimization of completion times or another time-based objective function.

The particular challenge of this problem variant arises from the interaction between the recharging decisions and the time-related constraints and objective. Any time spent for recharging (the detour to visit a station as well as the actual recharging duration) postpones the completion times of all jobs following a recharging stop, which generally favors short and late, i.e., with few succeeding jobs, recharges. The interaction is particularly pronounced if PR is applied, where also the recharging duration is part of the decision (see Fig. 1). When recharging, this results in a direct tradeoff between the battery level and the time: the longer the recharging duration the higher the battery level, i.e., the longer the vehicle range, but the later the time, i.e., the more unfavorable regarding the completion times, and vice versa.

1.1. Contributions

In this paper, we investigate the problem of route planning for a fleet of electric AMRs in intralogistics under different recharging strategies. The associated contributions are threefold.

First and to formalize the problem, we introduce a new family of EVRPs coined the electric vehicle routing problem with due dates minimizing completion times (EVRPD-C). By applying the three common recharging strategies BS, FR, and PR, we consider three different variants of the problem denoted EVRPD-C(BS), EVRPD-C(FR), and EVRPD-C(PR), respectively. The novelty here lies in the new completion-time based objective (CT) which is a main characteristic of the considered intralogistic AMR route planning problem. Due to the interplay of objective function and recharging decisions, the EVRPD-C is interesting not only from a practical but also from a methodological point of view. For comparison reasons, we also consider the corresponding problem EVRPD-D for all recharging strategies, which takes the standard VRP objective of minimizing the total traveled distance (DI) instead of the completion times.

Second, we propose an effective branch-price-and-cut (BPC) approach based on state-of-the art techniques for VRPs for solving the EVRPD-C. The core of the approach are ad-hoc labeling algorithms tailored to the specific variant (BS, FR, or PR) for solving the column-generation (CG) subproblems. The main novelty here relates to the EVRPD-C(PR), for which the pricing subproblem is considerably more complex compared to the equivalent problem minimizing distances. We introduce new resource extension functions (REFs) and dominance rules modeling the tradeoff between battery level and the time, which translates into costs in the case of minimizing completion times. For the EVRPD-C(BS) and EVRPD-C(FR), the labeling algorithms are very much related to their distance-minimizing counterparts that can be adapted straightforwardly. The same BPC is also used for solving the EVRPD-D variants employing slightly modified versions (see Section 4.1.1 for details) of the labeling algorithms of Desaulniers et al. (2016) for solving the subproblems.

Third, we report an extensive computational study generating managerial and computational insights based on two different sets of benchmark instances. The first is derived from the well-known Solomon instances and is dedicated to planning problems of electrically powered trucks (Schneider et al., 2014). These instances are mainly included to assess the computational performance of our algorithm on well-established instances. The second is a newly established benchmark set derived from a real-world intralogistics application. These instances feature the typical characteristics of AMRs in intralogistics which differ considerably from road transportation with electric trucks. We provide new managerial insights on the intralogistics AMR route planning problem with respect to the different recharging strategies and objective functions based on a comprehensive analysis of the problems' and solution's features. All instances and their corresponding solutions are publicly accessible on Zenodo (Meyer, Gschwind, Amberg, & Colling, 2024).

1.2. Outline

The remainder of this paper is organized as follows. In Section 2, we provide an overview of previous research on route planning for AMRs and briefly discuss exact approaches to VRP variants that are closely related to the EVRPD-C. Section 3 describes the setting of our intralogistics AMR route planning problem, formally defines the derived EVRP variants, and presents the extended set-partitioning formulation used for their solution. The details of our BPC approaches are provided in Section 4, with a focus on the labeling algorithms. In Section 5, we introduce the benchmark instances and present the results of our computational study. Section 6 summarizes the most important insights about the impact of the objective functions and recharging strategies. In Section 7, we conclude the paper with a summary and identify areas for future work.

2. Related work

The recent surveys (De Ryck, Versteyhe and Debrouwere, 2020; Fragapane et al., 2021; Le-Anh & De Koster, 2006; Vis, 2006) provide a good general overview of planning and control problems in the context of fleets of conventional autonomous vehicles, often referred to as autonomous guided vehicles or AGVs, and AMRs (note that while there exists no clear definition or distinction of AGV and AMR, the term AMR has become more common for freely navigating vehicles). However, these surveys only briefly touch on battery management. In contrast, solution approaches for variants of the EVRP were subject to a large number of publications in the last decade (Erdelić & Carić, 2019; Kucukoglu, Dewil, & Cattrysse, 2021). However in this area, most of the authors had road transportation with electric trucks or cars in mind, which has different characteristics than AMRs in an intralogistics setup (we refer to Section 5 for details on these differences). We, therefore, focus on reviewing approaches dedicated to route planning for autonomous vehicles in intralogistics (Section 2.1) that also take battery management into account or have a time-based objective function. Furthermore, in Section 2.2, we briefly review previous works on exact approaches to closely related VRPs. Note that in the intralogistics context, transport *tasks* typically are of pickup-and-delivery type. In EVRPs, on the other hand, the *tasks* to be performed are of single-visit nature. Therefore, we speak of (transport) *jobs* in the former case and *customers* in the latter.

2.1. AMRs in intralogistics systems

In practice, rule-based approaches are often applied to incorporate recharging decisions into vehicle routing in the intralogistics context. There exist a couple of studies investigating the impact of these rules on the system's performance (e.g. Ebben, 2001; Kabir & Suzuki, 2018, 2019; McHaney, 1995; Zou et al., 2018). In general, they differentiate between opportunity recharging, automatic recharging, and a combination of both (e.g. De Ryck, Versteyhe, Debrouwere, 2020; De Ryck, Versteyhe, Shariatmadar, 2020; McHaney, 1995). In opportunity recharging, the vehicles use idle times for recharging. In automatic recharging, the vehicles drive to a recharging station if the battery is depleted to a certain threshold and fully recharge the battery or recharge it to a given threshold. The recharging station is selected by rules, such as the closest station or the first station on the current route. The studies show that recharging decisions can have a considerable impact on the system's performance. This is especially true if no natural breaks, such as off-shift times, can be used for recharging and if the vehicles' share of idle times does not exceed 50% (Kabir & Suzuki, 2018; Mangiaracina, Perego, Seghezzi, & Tumino, 2019).

To the best of our knowledge, there are only a few more sophisticated optimization approaches to incorporate recharging decisions into vehicle routing in the intralogistics context. An auction-based approach for assigning vehicles to recharging stations is proposed in Selmair, Hauers, and Gustafsson-Ende (2019). The objective of the auction-mechanism considers the distance, occupation cost, and cost for the difference of battery levels among the competing vehicles. The auction-mechanism is applied as part of an opportunity recharging strategy. Studies using a discrete-event simulation model show that frequent and fast opportunity recharging is advantageous compared to infrequent but long recharging because it can increase the number of jobs that can be performed.

De Ryck, Versteyhe, Shariatmadar (2020) propose a particle swarm optimization algorithm to add recharging stations to a predefined route of a vehicle. The algorithm determines which recharging station should be included at which position in the route and how long the vehicle should recharge according to a partial (opportunity) recharging strategy. The objective is to minimize the total traveling time including the time for recharging. The approach is proposed as a base for decentralized resource management. Computational experiments show that the total traveling time can be decreased compared to an automatic recharging strategy.

Jun et al. (2021) addressed effective route planning of a fleet of electric AMRs under application-specific conditions. They point out that the battery capacity of AMRs is limited and that the main objective is to respect the due dates of the transport requests rather than the minimization of travel cost. The authors propose a mixed integer linear program (MIP) for the pickup and delivery problem with due dates and partial recharges for minimizing the total tardiness that can solve small instances to optimality. For larger instances with up to 50 jobs and ten AMRs, the authors proposed two simple heuristics for a construction phase and three heuristics for an improvement phase. In contrast to the MIP, the heuristic approaches only support a full recharge strategy. Their results show that the total tardiness can be reduced by up to one quarter if recharging with partial recharges is applied on instances with a small number of transport requests and AMRs that are initially low in battery. This effect decreases with increasing instance size but is still clearly present in some cases.

The various studies show that the choice of a suitable recharging strategy depends on the problem under consideration. If flexibility in recharging is possible, opportunity recharging seems beneficial to automatic recharging.

2.2. Exact algorithms for related VRPs

The main characteristics of the EVRPD-C are the inclusion of recharging decisions and an objective that minimizes completion times. In the following, we briefly discuss the main exact approaches for closely related vehicle routing problems.

In Desaulniers et al. (2016), the first exact solution method for the EVRP with time windows (EVRPTW) is introduced. The authors propose BPC algorithms for four different variants of the EVRPTW with regard to the recharging strategy: single and multiple recharges allowed per route as well as full and partial recharges. For all variants, the objective is to minimize total travel costs. The recharging time is assumed to be linear. To evaluate the approach, instances are generated based on the EVRPTW instances of Schneider et al. (2014) (which are based on the 100-customer benchmark instances of Solomon (1987) for the VRPTW). Instances with up to 100 customers and 21 recharging stations can be solved to optimality. The authors point out that routing costs and the number of employed vehicles can be reduced when multiple and partial recharges are allowed (by 2.6% and by 4.0% on average compared to full recharges). Then, the number of recharges per vehicle increases significantly (by 20.7% on average). Duman, Taş, and Çatay (2022) adopt the approach of Desaulniers et al. (2016) to develop an exact BPC algorithm and a BPC-based heuristic for the EVRPTW variants with partial recharges. They are able to provide a few new best known solutions. Finally, Desaulniers, Gschwind, and Irnich (2020) propose an acceleration strategy for BPC algorithms on path-based models, including the four EVRPTW variants considered in Desaulniers et al. (2016).

Their results show that the acceleration strategy can reduce overall computation times, and many previously unsolved instances can be solved to optimality.

In the context of EVRPs, Montoya, Guéret, Mendoza, and Villegas (2017) introduced the application of non-linear recharging functions, i.e., piecewise-linear approximations of the real recharging function. Schulz (2024) proposes a first exact solution approach for an EVRPTW with a non-linear and concave recharging function (EVRPTW-NL). In contrast to former works, his piecewise linearization of the recharging function never underestimates the real battery level. He develops a branch-and-cut approach that ensures an optimal solution for the real recharging function. A computational study on generated instances with 10–100 customers and 2–10 recharging stations shows that the average number of customers per vehicle mainly determines whether an instance is difficult to solve. The number of recharging stations and the length of the time windows, on the other hand, have less influence.

Keskin, Laporte, and Çatay (2019) assume that recharging stations (in road transportation networks) have a limited capacity so that arriving vehicles have to wait if the station is occupied. The authors model an EVRPTW with a piecewise-linear recharging function and time-dependent waiting times at the stations (EVRPTW-NL-WS). They formulate the problem as a mixed integer linear program and propose a matheuristic that combines an Adaptive Large Neighborhood Search with the solution of the MILP. To evaluate their approach, the authors adapt the instances of Schneider et al. (2014). They narrow the time windows to increase their impact on recharging decisions. Small instances with 5–10 customers can be solved with a general-purpose MIP solver. These solutions are used to evaluate the performance of the heuristic approach. In general, the results indicate that waiting times can have an impact on routing decisions.

Wang, Adulyasak, Cordeau, and He (2024) investigate routing a heterogeneous fleet with multiple recharging modes and time-dependent waiting times at recharging stations (WS). Their EVRPTW-NL-WS is modeled as mixed integer linear programming formulation and evaluated on modified instances of Hiermann, Puchinger, Ropke, and Hartl (2016) and instances of Montoya et al. (2017). For instances with up to 320 customers, the problem is solved heuristically with an iterated local search algorithm and a set partitioning model. For small instances with 5–15 customers and 4–6 different vehicle types, their formulation can be solved with a general-purpose MIP solver. However, the small instances are used to evaluate the heuristic's performance and not to draw conclusions about the consideration of the various problem characteristics.

Lam, Desaulniers, and Stuckey (2022) combine the examination of non-linear recharging functions (NL) and capacitated recharging stations (CS) in the context of EVRPTWs. Their EVRPTW-NL-CS contains a piecewise-linear recharging function and a scheduling part to model recharging around the availability of recharging stations. They propose a hybrid BPC algorithm where the scheduling part is solved with constraint programming using logic-based Benders decomposition. The authors evaluate their approach on modified instances of Schneider et al. (2014) comprising 25–100 customers. The modified instances contain fewer recharging possibilities, and the capacities of the recharging stations are limited to one or two vehicles. Their results indicate that the capacity restrictions of the stations are not crucial in most cases. They also show that an additional recharging capacity at each station significantly reduces the capacity restrictions.

Lera-Romero, Bront, and Soulignac (2024) study the effects of time-dependent travel times and energy consumption in transportation networks such as last-mile distribution in large cities and mid-haul logistics in retail. They propose a formulation for the time-dependent EVRPTW that includes time-dependent battery consumption, non-linear recharging times, and time-dependent waiting times at recharging stations (TD-EVRPTW-NL-WS). The authors develop a BPC framework that extends and generalizes the work of Desaulniers et al. (2016). For their experiments, they extend the instances with 25–100 customers from Desaulniers et al. (2016) with time-dependent information, non-linear recharging functions, and time-dependent waiting time functions. They evaluate solutions obtained from a time-independent model with the time-dependent model and report average time window and battery capacity violations of 1.4 to 3.6% and 0.5 to 1.1%, respectively.

Gschwind, Irnich, Tilk, and Emde (2019) investigate a variant of the VRP with time windows minimizing weighted completion times that arises when planning direct delivery trips in a just-in-time context. Recharging decisions are not part of the model. The proposed BPC algorithm consistently solves instances with 25 to 200 delivery trips. For an overview of related problems minimizing completion times, such as parallel machine scheduling, we refer to the literature presented in Gschwind et al. (2019).

Table 1 provides an overview of the most important characteristics of closely related problems in order to show the differences to the EVRPD-C considered here. For each problem type (*Type*), we summarize main references to exact solution approaches (*Main References*), objectives and main constraints (*Obj. (Main Const.)*), battery recharging strategies and assumptions (*Recharging*), and characteristics of the considered fleet and optimally solved instance sizes (*Fleet and Inst.*).

Most EVRPTW variants are evaluated on the instances of Desaulniers et al. (2016) and Schneider et al. (2014) that are dedicated to road transportation with electric trucks or cars. This setting is also the basis for most of the other instance sets. With regard to the choice of a suitable recharging strategy, it can be summarized that for these benchmark instances, multiple recharges per route are more beneficial than single recharges, and PR is more beneficial than FR. A combination of multiple recharges and PR usually leads to the greatest savings in terms of routing costs and the number of vehicles used. Our findings on the impact of recharging strategies for AMR route planning in an industrial context are presented in Section 6.

3. Problem description and mathematical formulation

In this section, we first present the background and the assumptions of the intralogistics AMR route planning problem. We then describe the EVRPD-C, which formalizes our practical planning problem, and the considered strategies BS, FR, and PR. Finally, an extended set-partitioning formulation of the problem is given.

3.1. AMR route planning problem in intralogistics

In intralogistics, AMRs are used, e.g., to supply the assembly stations of a production line or of a workshop production. Our focus is on the route planning for a fleet of electric AMRs, including battery management, given a set of transport jobs to be performed. In the following, we describe the practical problem setup in more detail.

AMR characteristics. In operative decision making, a usually fixed fleet of AMRs is available with known characteristics for each vehicle, such as battery capacity, speed with and without load, or battery consumption rates for driving and material handling.

- *Capacity:* AMRs often have a capacity of one unit, as they are able to transport a single pallet, rack, or small load carrier at a time.
- *Range and recharging:* Each vehicle has a given battery level at the beginning and a minimum battery threshold to meet at the end of its route. The driving range of AMRs is limited so that several dedicated stops per shift or day are necessary if the AMRs do not have a large share of idle times. State-of-the-art AMRs are able to work autonomously with BS or IR. For induction-based recharging, FR and PR strategies can be applied.
- *Positioning:* Typically, the AMRs are not based at a common depot but are spread over the production site. Hence, each vehicle has a different start and end location for the route planning. Obviously, the final battery threshold should be sufficient to drive to the next recharging station.

Transport jobs. In many applications, the material handling part is fully automated, and vehicles must not wait in front of a station to take over or transfer a transport unit. The transport jobs are pickup and delivery jobs, each with individual pickup and delivery locations. Each job has a due date specifying the latest possible time at which the transport unit is to be delivered at the delivery location. In some cases, jobs also have an availability time for the pickup which we do not consider in the following.

Table 1
Overview of the most important characteristics of related problems and main references to exact solution approaches. *NL*: non-linear recharging, *WS*: time-dependent waiting times at recharging stations, *CS*: capacitated recharging stations, *TD*: time-dependent, *PD*: pickup and delivery, *D*: due dates, *VEH*: number of vehicles, *DI*: distance (or travel time), *TR*: time for recharging, *TV*: time window violations, *TT*: total tardiness, *CT*: sum of completion times, *TW*: time windows, *SW*: soft time windows, *LC*: load capacity, *BC*: battery capacity, *SR*: single recharge per route, *MR*: multiple recharges per route, *BS*: battery swapping, *FR*: full-recharges strategy, *PR*: partial-recharges strategy, *HO*: homogeneous vehicle types, *HE*: heterogeneous vehicle types, *SD*: single depot, *MD*: multiple depots, *?*: not reported.

Type	Main references	Obj. (Main Const.)	Recharging	Fleet and Inst.
VRPTW	Gschwind et al. (2019)	min VEH, CT (TW)	No recharging	HO, SD 25–200 cust.
EVRPTW	Desaulniers et al. (2016)	min DI (TW, LC, BC)	SR/MR: FR, PR (linear)	HO, SD 25–100 cust.
EVRPTW-NL	Schulz (2024)	min DI, TR (TW, BC)	MR: PR (non-linear)	HO, SD 10-100 cust.
EVRPTW-NL-WS	Keskin et al. (2019)	min VEH, DI, TV (SW, LC, BC)	MR: PR (non-linear)	HO, SD 5–10 cust.
	Wang et al. (2024)	min VEH, DI, TR (TW, LC, BC)	MR: PR (non-linear)	HE, SD 5–15 cust.
EVRPTW-NL-CS	Lam et al. (2022)	min DI (TW, LC, BC, CS)	MR: FR, PR (non-linear)	HO, SD 25-? cust.
TD-EVRPTW-NL-WS	Lera-Romero et al. (2024)	min DI (TW, LC, BC)	SR/MR: FR, PR (non-linear)	HO, SD 25–100 cust.
EVRPD-PD	Jun et al. (2021)	min TT (D, LC, BC)	MR: BS, FR, PR (linear)	HO, SD 5–10 jobs
EVRPD-D	This work	min DI (D, LC, BC)	MR: BS, FR, PR (linear)	HO, MD 12–96 cust.
EVRPD-C	This work	min CT (D, LC, BC)	MR: BS, FR, PR (linear)	HO, MD 12–108 cust.

Objectives. As the traveled distances in intralogistics applications are comparatively small and idle production resources are very expensive, the main objective is typically either the minimization of the total tardiness or the minimization of the sum of completion times. Our focus is on the latter variant. By finishing jobs early, the risk of delays and disruptions can be reduced irrespective of the due date.

Further assumptions. Apart from the characteristics introduced so far, we make the following assumptions:

- The recharging process is modeled in a linear fashion. This is a common simplification for planning algorithms (e.g., Schneider et al., 2014). To use non-linear recharging functions, in literature, often piecewise linear approximations of the recharging functions are applied (e.g., Montoya et al., 2017). For a recent overview on approaches with linear and non-linear recharging functions, we refer to Schulz (2024).
- The capacity of recharging stations in terms of the number of parallel recharges is not restricted (no waiting). This assumption typically holds if there is one recharging station with a large number of parallel recharging mats or battery swapping appliances, or if recharging stations with a capacity larger than one are well distributed over the site.
- Vehicles are available during the whole planning period. In practical applications, this assumption is largely valid, unless vehicles will be withdrawn for maintenance or inspection work or other tasks.
- Vehicles only recharge between two jobs, i.e., without any load. This restriction corresponds to the procedure in practice.

3.2. Electric vehicle routing problem with due dates

According to the description in the previous section, the considered intralogistics route planning problem is a variant of an EVRP with pickup and delivery. However, as AMRs have a unit load capacity and can only recharge between two transport jobs, the problem can equivalently be modeled as an EVRP in which pickup and delivery jobs are represented by visits to a single *customer* and vehicles are empty before/after each of these visits. Recall that we refer to *customers* in the general context of EVRPs while we use the term *jobs* when addressing specifically the intralogistics context. For conciseness reasons, we omit the detailed description of the corresponding pickup and delivery EVRP and directly present the EVRPD-C in the following. The problem transformation is described in detail in Appendix A.

The EVRPD-C can be defined on a directed graph $G = (V, A)$. The vertex set $V = N \cup R \cup O \cup D$ comprises the set of customers N, the set of recharging stations R, and the sets of origin and destination vertices O and D, respectively. The arc set is $A = \{(i,j) \in V \times V | i \neq j\}$.

Let K be the set of different vehicle types. A vehicle type $k \in K$ is characterized by origin vertex $o_k \in O$, destination vertex $d_k \in D$, load capacity Q_k, initial battery level H_k^{init}, battery capacity H_k^{max}, and final battery threshold H_k^{end}. Two vehicles are of the same type if they have the same characteristics. The number of vehicles of type k is m_k. The set of vertices and arcs relevant for vehicle type k are $V_k = N \cup R \cup \{o_k\} \cup \{d_k\}$ and $A_k = \{(i,j) \in A | i, j \in V_k\}$, respectively.

With each vertex $i \in V$ is associated a latest possible start of service l_i, corresponding to the due dates for the service at customer vertices, while we assume l_i to be not binding for all recharging stations and origin/destination vertices. Furthermore, a non-negative demand $q_i \geq 0$ is associated with each vertex i, with $q_i = 0$ for $i \in V \setminus N$. Recall that in the AMR route planning context, customer demands and vehicle capacities are not involved. As we also solve benchmark instances from the literature that include customer demands and vehicle capacities, we keep them in the problem and algorithm descriptions.

A travel time t_{ij}, a travel distance c_{ij}, and a battery consumption h_{ij}, all of which are assumed to be non-negative and satisfy the triangle inequality, are associated with each arc $(i,j) \in A$. The travel times t_{ij} include a possible service duration at vertex i. Note that we assume an identical speed profile and battery consumption for all vehicle types. The generalization of problem description and solution approach to vehicle-type specific values is straightforward.

As proposed by Desaulniers et al. (2016), all battery related parameters (H_k^{init}, H_k^{max}, H_k^{end}, h_{ij}) are given in recharging time units, e.g., h_{ij} is the time it takes to recharge the amount of energy that is consumed for traveling from vertex i to j. Consequently, there is a one-to-one relation between recharging duration and the battery's state of charge in IR, i.e., recharging for Δ units of time increases the battery level by exactly Δ units. In case of BS, the duration of a recharging stop is b. The conversion of energy in recharging time units simplifies the description and formulas and avoids numerical rounding errors in the computational results.

The EVRPD-C consists of finding a set of feasible vehicle routes, at most m_k for each vehicle type $k \in K$, such that each customer is visited exactly once and the sum of the starts of service at the customer nodes is minimal (which is equivalent to the minimization of the sum of completion times). If necessary, a vehicle can visit one or several (not necessarily different) recharging stations. There is no limit on the number of visits to recharging stations, neither for an individual vehicle or recharging station nor in total. Recharging is performed according to one of the recharging strategies BS, FR, or PR giving rise to three different variants of the EVRPD-C denoted by EVRPD-C(BS), EVRPD-C(FR), and EVRPD-C(PR), respectively. With BS, the vehicle battery is swapped with a fully recharged one at each recharging stop. The recharging duration equals the constant b. With FR and PR, vehicles are recharged by induction, and the amount of energy recharged depends on the recharging duration. With FR, the vehicles have to be recharge completely at each recharging stop, while with PR any amount of energy can be recharged.

A feasible route for vehicle type $k \in K$ is an elementary o_k-d_k-path in $G_k = (V_k, A_k)$ that respects the vehicle capacity Q_k and for which a feasible pair of schedule and battery-recharging plan exists. Let an elementary o_k-d_k-path $(o_k = i_0, i_1, \dots, i_p = d_k)$ for some $p > 1$ be given. A pair of schedule $(T_0, T_1, \dots, T_p)$ and battery-recharging plan $(X_0, X_1, \dots, X_p)$ is feasible is it follows the rules of the given recharging strategy and fulfills the following conditions:

(i) recharging is only allowed at recharging stations
(ii) the service time limits l_i at all visited vertices i are met
(iii) the battery charge level is never negative and never above the battery capacity H_k^{max}
(iv) when arriving at the destination d_k the battery charge level is not smaller than the threshold H_k^{end}

We now formalize these conditions for the three considered recharging strategies. For FR and PR, condition (i) translates to $X_j = 0$ for $i_j \in V\backslash R$ and $0 \le X_j \le H_k^{max}$ for $i_j \in R$. Condition (ii) requires $T_{j-1} + X_{j-1} + t_{i_{j-1},i_j} \le T_j$ and $T_j \le l_j$ to hold for all $j = 1, \dots, p$. The battery limits of condition (iii) imply $H_k^{init} - \sum_{j=1}^{q} h_{i_{j-1},i_j} + \sum_{j=1}^{q-1} X_j \ge 0$ and $H_k^{init} - \sum_{j=1}^{q} h_{i_{j-1},i_j} + \sum_{j=1}^{q} X_j \le H_k^{max}$ for all $q = 1, \dots, p$. Finally, the final battery threshold of condition (iv) is given by $H_k^{init} - \sum_{j=1}^{p} h_{i_{j-1},i_j} + \sum_{j=1}^{p-1} X_j \ge H_k^{end}$. For FR, the following additional condition requires the battery to be fully recharged at all recharging stations: $H_k^{init} - \sum_{j=1}^{q} h_{i_{j-1},i_j} + \sum_{j=1}^{q} X_j = H_k^{max}$ if $i_q \in R$.

For strategy BS, only some slight modifications in these formalizations are necessary. First, the duration of all stops at recharging stations now equals the constant b, independent of the amount of energy that is restored by the swap. We, therefore, modify the service time propagation of condition (ii) to become $T_{j-1} + b + t_{i_{j-1},i_j} \le T_j$ if $i_q \in R$. Second, the same condition as for FR can model the fact that after each stop at a recharging station, the vehicle is equipped with a fully recharged battery.

For comparison reasons, we also consider the related problem EVRPD-D for all three recharging strategies. The only difference between EVRPD-C and EVRPD-D is that the latter uses the standard VRP objective of minimizing the total traveled distances. Feasibility of a route is not influenced by the altered objective function.

3.3. Mathematical formulation

Let Ω_k denote the set of feasible routes for a vehicle of type $k \in K$. Depending on the selected objective, each route p is associated with a cost c_p consisting either of the sum of service start times of the customers visited on p (EVRPD-C) or of the total traveled distance of route p (EVRPD-D). Parameter a_{pn} with $n \in N$ is a binary parameter indicating if customer n is visited by route p. Binary variable θ_p becomes one if route p is part of the solution and 0 otherwise.

Based on the notation above, we can formulate all problem variants of the EVRPD as follows:

$$\min \sum_{k \in K} \sum_{p \in \Omega_k} c_p \theta_p, \tag{1a}$$

$$\text{s.t.} \sum_{k \in K} \sum_{p \in \Omega_k} a_{pn} \theta_p = 1 \quad n \in N, \tag{1b}$$

$$\sum_{p \in \Omega_k} \theta_p \le m_k \quad k \in K, \tag{1c}$$

$$\theta_p \in \{0, 1\} \quad k \in K, p \in \Omega_k. \tag{1d}$$

The objective function (1a) minimizes the total cost while partitioning constraints (1b) assure that each customer is visited exactly once. Constraints (1c) limit the number of vehicles of each type. Constraints (1d) define the domain of the decision variable.

4. Branch-price-and-cut algorithms

As the number of feasible routes is typically too large, formulation (1) cannot be solved directly, and we employ a BPC algorithm for its solution. A BPC algorithm is a branch-and-bound algorithm in which the lower bounds are computed using CG. Cuts are added dynamically to strengthen the linear relaxations. CG is an iterative procedure that alternates between the solution of a restricted master problem (RMP) and the solution of one or several pricing problems. Our RMP is the linear relaxation of model (1) in which the sets of feasible routes Ω_k are replaced by proper subsets $\Omega'_k \subset \Omega_k, k \in K$. The task of the pricing problems (Section 4.1), one for each vehicle type $k \in K$, is to dynamically identify routes (=columns) with negative reduced cost, if any exist, which are then added to the RMP. If no negative reduced-cost routes exist, the current linear relaxation is solved to optimality. The corresponding lower bound can be strengthened by adding valid inequalities (Section 4.2), and branching is required to guarantee integer solutions (Section 4.3).

4.1. Pricing problems

Let $\pi_n, n \in N$ and $\mu_k, k \in K$ be the dual variables corresponding to constraints (1b) and (1c), respectively, with $\pi_n \in \mathbb{R}$ and $\mu_k \le 0$. For all problem variants of the EVRPD, the reduced cost of a route $p \in \Omega_k$ is given by $\tilde{c}_p = c_p - \sum_{n \in N} a_{pn}\pi_n - \mu_k$ and the pricing problem for vehicle type $k \in K$ is given by

$$\min_{p \in \Omega_k} \tilde{c}_p. \tag{2}$$

Similar to many other VRPs, the pricing problems (2) are variants of elementary shortest path problems with resource constraints depending on the considered EVRPD variant. They can be solved by means of dynamic-programming labeling algorithms (Irnich & Desaulniers, 2005). In a labeling algorithm, partial paths are iteratively extended looking for a minimum-cost path from a given source to a given sink. The partial paths are implicitly represented by labels that store the accumulated reduced cost and resource consumption along the partial paths. The labels are propagated along the network arcs by resource extension functions (REFs). To avoid the enumeration of all feasible paths, dominance rules are employed to eliminate labels that cannot lead to an optimal solution of (2). We refer to Irnich and Desaulniers (2005) for a more comprehensive description of labeling algorithms.

4.1.1. Labeling algorithms for EVRPD-D variants

For the variants with the objective of minimizing the total distance of vehicle routes, we define $\tilde{c}_{ij}^k = c_{ij} - 1/2\tilde{\pi}_i - 1/2\tilde{\pi}_j$ with $\tilde{\pi}_n = \pi_n$ for all customers $n \in N$, $\tilde{\pi}_r = 0$ for all recharging stations $r \in R$, and $\tilde{\pi}_{o_k} = \tilde{\pi}_{d_k} = \mu_k$. With these definitions, the reduced cost $\tilde{c}_p$ of a route $p \in \Omega_k$ is $\tilde{c}_p = \sum_{(i,j) \in A(p)} \tilde{c}_{ij}^k$, where $A(p)$ denotes the sequence of arcs traversed by route p.

Regarding the FR and PR cases, Desaulniers et al. (2016) have derived ad-hoc labeling algorithms for the more general problem variants with time windows, i.e., with lower and upper limits for the start of service. For the sake of completeness, we also present the core components of the labeling algorithms for these variants, which are simplified versions of those presented in Desaulniers et al. (2016), enabled by the absence of lower time window limits in our case. Please note that the resources we use slightly differ from those in Desaulniers et al. (2016), but there is a direct relation between them. We feel, however, that our resources are more intuitive to interpret. Furthermore, we present details of the labeling algorithm for the BS variant, which results from a rather straightforward adaptation of the FR case.

EVRPD-D with full recharges. For the EVRPD-D(FR), a partial path p of a vehicle of type $k \in K$ from its origin o_k to a vertex $i \in V$ is represented by a label $L_i = (T_i^{cost}, T_i^{load}, T_i^{tMin}, T_i^{rt}, (T_i^{unreach_n})_{n\in N})$. The components of the label have the following meaning:

- T_i^{cost} the reduced cost of path p;
- T_i^{load} the total load delivered along path p;
- T_i^{tMin} the earliest start of service at vertex i;
- T_i^{rt} the cumulated required recharging time since the last recharge (if any) or since leaving the depot (otherwise). It is used to assure that the maximum battery capacity is not exceeded, and, in case of FR, it determines the recharging time in recharging stations.
- $T_i^{unreach_n}$ binary parameters indicating if a customer $n \in N$ is reachable ($T_i^{unreach_n} = 0$) or unreachable ($T_i^{unreach_n} = 1$). A customer n is unreachable if it has already been visited on p or if traveling directly from i to n violates its latest possible service time l_n or the vehicle capacity Q_k, otherwise it is reachable.

The initial label at the origin o_k is given by $L_{o_k} = (0, 0, 0, H_k^{max} - H_k^{init}, \mathbf{0})$. The extension of a label L_i along an arc $(i,j) \in A_k$ is feasible if the following conditions hold:

$$T_i^{load} + q_j \le Q_k, \tag{3a}$$

$$T_i^{tMin} + t_{ij} \le l_j, \tag{3b}$$

$$T_i^{rt} + h_{ij} \le \begin{cases} H_k^{max} - H_k^{end} & \text{if } j = d_k, \\ H_k^{max} & \text{otherwise,} \end{cases} \tag{3c}$$

$$T_i^{unreach_j} = 0 \quad \text{if } j \in N. \tag{3d}$$

If the extension is feasible, a new label $L_j = (T_j^{cost}, T_j^{load}, T_j^{tMin}, T_j^{rt}, (T_j^{unreach_n})_{n\in N})$ is created according to the following REFs:

$$T_j^{cost} = T_i^{cost} + \tilde{c}_{ij}^k, \tag{4a}$$

$$T_j^{load} = T_i^{load} + q_j, \tag{4b}$$

$$T_j^{tMin} = \begin{cases} T_i^{tMin} + t_{ij} + T_i^{rt} & \text{if } i \in R, \\ T_i^{tMin} + t_{ij} & \text{else,} \end{cases} \tag{4c}$$

$$T_j^{rt} = \begin{cases} h_{ij} & \text{if } i \in R, \\ T_i^{rt} + h_{ij} & \text{otherwise,} \end{cases} \tag{4d}$$

$$T_j^{unreach_n} = \begin{cases} 1 & \text{if } j = n \ \vee \ T_j^{load} + q_n > Q_k \ \vee \ T_j^{tMin} + t_{jn} > l_n, \\ T_i^{unreach_n} & \text{otherwise.} \end{cases} \tag{4e}$$

Since the REFs (4) are non-decreasing (see, e.g., Irnich, 2007) and the label components are bounded only from above by (3), the following dominance rule to eliminate unpromising labels is directly valid (Irnich & Desaulniers, 2005). Let $L^k = (T_k^{cost}, T_k^{load}, T_k^{tMin}, T_k^{rt}, (T_k^{unreach_n})_{n\in N})$, $k \in \{1,2\}$, be two labels associated with different paths ending at the same vertex. Label L^1 dominates label L^2 if

$$T_1^{cost} \le T_2^{cost}, \tag{5a}$$

$$T_1^{load} \le T_2^{load}, \tag{5b}$$

$$T_1^{tMin} \le T_2^{tMin}, \tag{5c}$$

$$T_1^{rt} \le T_2^{rt}, \tag{5d}$$

$$T_1^{unreach_n} \le T_2^{unreach_n} \quad n \in N. \tag{5e}$$

Note that in case of zero demands for all customers, the load component T_i^{load} together with the associated feasibility condition (3a), REF (4b), and dominance condition (5b) can be omitted. The same holds true for all other problem variants considered in the following.

EVRPD-D with battery swapping. The labeling algorithm for the EVRPD-D(BS) is almost identical to the FR case. In fact, the same label components, initial label, the feasibility conditions (3), and the dominance rule (5) of the FR case are also used in the BS case. The only difference is in the REFs (4), where we need to substitute (4c) by

$$T_j^{tMin} = \begin{cases} T_i^{tMin} + t_{ij} + b & \text{if } i \in R, \\ T_i^{tMin} + t_{ij} & \text{otherwise.} \end{cases} \tag{6}$$

EVRPD-D with partial recharges. With an PR strategy, the amount to recharge (=the recharging duration) at a recharging station can be chosen freely. In general, there is a tradeoff between the battery level (implying the potential range) and the time: the longer the recharging duration, the higher the battery level, i.e., the longer the range, but the later the time, i.e., the more unfavorable regarding the customers' due dates, and vice versa. Even more, the amount needed to recharge at a recharging station depends on the route-part after visiting this recharging station, which is unknown for a partial path in a labeling algorithm. Therefore, additional resources are needed to express and propagate the corresponding tradeoff function.

A Label L_i representing a partial path p of a vehicle of type $k \in K$ from its origin o_k to a vertex $i \in V$ comprises the following resources $L_i = (T_i^{cost}, T_i^{load}, T_i^{tMin}, (T_i^{unreach_n})_{n\in N}, T_i^{nRch}, T_i^{remBat}, T_i^{potRange})$, where T_i^{tMin} has a slightly modified meaning compared to the FR and BS cases. The modified and additional components are as follows:

- T_i^{tMin} the earliest start of service at vertex i. If a recharging station is visited on p prior to i, T_i^{tMin} incorporates a minimum recharge, if necessary, to reach vertex i battery feasible. If no recharging station is visited prior to i, T_i^{tMin} is the standard earliest start of service.
- T_i^{nRch} the number of recharging stations visited along path p;
- T_i^{remBat} the remaining battery level (from the initial level) at vertex i;
- $T_i^{potRange}$ the maximum battery level which can be reached if all the time available for recharging is used for recharging at the preceding recharging station, taking into account the maximum battery level H_k^{max} and the due dates of all customers between the recharging station and vertex i. It corresponds to the potential range of the vehicle. If no recharging station is visited on p prior to i, $T_i^{potRange}$ equals the remaining battery level T_i^{remBat}. In $T_i^{potRange}$ only the potential recharging time of the last visited recharging station must be considered: An increase in recharging time at a prior recharging station would result in a higher remaining battery level on arrival at the last visited recharging station. Hence, the potential of the last visited recharging station would be reduced by the increase in recharging time at the prior recharging station.

Resources T_i^{tMin}, T_i^{remBat}, and $T_i^{potRange}$ model the described tradeoff between battery level and time as depicted in Fig. 2 (left plot). The tradeoff function represents the possible battery level when reaching vertex i given a specific start of service at vertex i. T_i^{tMin} is the earliest time vertex i can be reached, the corresponding battery level is T_i^{remBat}. It represents the remaining battery level (from the initial level), i.e., if $T_i^{remBat} > 0$ then no recharging is performed up to vertex i. Otherwise, i.e., if $T_i^{remBat} = 0$, then T_i^{tMin} includes the minimum duration for recharging that is necessary to feasibly reach vertex i (implying that

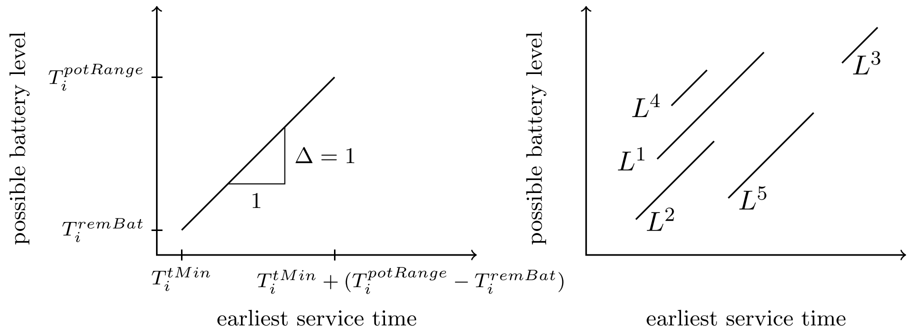

Fig. 2. Tradeoff between the earliest start of service time and the possible battery level.

a recharging station is visited on p prior to i). If a longer recharging duration is chosen at the recharging station preceding i on p, the battery level when reaching i increases but so does the start of service. Note that the tradeoff functions are linear with slope 1 due to the definition of the battery related parameters (see Section 3.2). The maximum possible battery level when reaching vertex i is $T_i^{potRange}$. It is constrained by either the maximum battery capacity or the due date of any customer visited after the recharging station preceding vertex i. The corresponding start of service at i can be computed as $T_i^{tMin} + (T_i^{potRange} - T_i^{remBat})$ and the maximum possible recharging duration at the station preceding i is $T_i^{potRange} - T_i^{remBat}$. If there is no recharging station visited on p prior to i, $T_i^{potRange}$ and T_i^{remBat} coincide and the tradeoff function diminishes to a single point.

The initial label at the origin o_k is $L_{o_k} = (0, 0, 0, \mathbf{0}, 0, H_v^{init}, H_v^{init})$. Extending a label L_i along an arc $(i, j) \in A_k$ is feasible if (3a), (3b), (3d), and

$$T_i^{potRange} \geq \begin{cases} h_{ij} + H_k^{end} & \text{if } j = d_k, \\ h_{ij} & \text{otherwise,} \end{cases} \tag{7}$$

hold. In case of feasibility, the label components are extended using REFs (4a), (4b), (4e), and the following components that model the propagation of the function for the tradeoff between battery level and (earliest service) time:

$$T_j^{nRch} = \begin{cases} T_i^{nRch} + 1 & \text{if } j \in R, \\ T_i^{nRch} & \text{otherwise,} \end{cases} \tag{8a}$$

$$T_j^{tMin} = \begin{cases} T_i^{tMin} + t_{ij} & \text{if } T_i^{nRch} = 0, \\ T_i^{tMin} + t_{ij} + \max\{0, h_{ij} - T_i^{remBat}\} & \text{otherwise,} \end{cases} \tag{8b}$$

$$T_j^{remBat} = \begin{cases} T_i^{remBat} - h_{ij} & \text{if } T_i^{nRch} = 0, \\ \max\{0, T_i^{remBat} - h_{ij}\} & \text{otherwise,} \end{cases} \tag{8c}$$

$$T_j^{potRange} = \begin{cases} H_k^{max} & \text{if } j \in R, \\ T_j^{remBat} & \text{if } j \notin R \wedge T_i^{nRch} = 0, \\ T_i^{potRange} - h_{ij} - \\ \max\{0, T_i^{potRange} - T_i^{remBat} \\ \quad -(l_j - T_i^{tMin} - t_{ij})\} & \text{otherwise.} \end{cases} \tag{8d}$$

The potential range $T_j^{potRange}$ is set to H_k^{max} whenever a recharging station is visited as we assume no time limits on recharging stations so that the recharging duration is not restricted. Along the path, $T_j^{potRange}$ is reduced by the consumption h_{ij} and possible restrictions on the recharging time at the preceding recharging station. The restrictions can result from the remaining battery level and the due date of customer j (see function (8d)).

To eliminate provably non-optimal partial paths, the following dominance rule can be applied. Let $L^k = (T_k^{cost}, T_k^{load}, T_k^{tMin}, (T_k^{unreach_n})_{n \in N}, T_k^{nRch}, T_k^{remBat}, T_k^{potRange})$, $k \in \{1, 2\}$, be two labels associated with different paths ending at the same vertex. Label L^1 dominates label L^2 if

$$T_1^{cost} \leq T_2^{cost}, \tag{9a}$$

$$T_1^{load} \leq T_2^{load}, \tag{9b}$$

$$T_1^{tMin} \leq T_2^{tMin}, \tag{9c}$$

$$T_1^{potRange} \geq T_2^{potRange}, \tag{9d}$$

$$T_1^{remBat} + T_2^{tMin} - T_1^{tMin} \geq T_2^{remBat}, \tag{9e}$$

$$T_1^{unreach_n} \leq T_2^{unreach_n} \quad n \in N. \tag{9f}$$

Conditions (9a), (9b), and (9f) are straightforward due to non-decreasing REFs and resources being bounded only from above, analog to the FR case. Conditions (9c)–(9e) model the dominance of L_i^1's tradeoff function over L_i^2's. More precisely, they ensure that for each relevant start of service $T \in [T_2^{tMin}, T_2^{tMin} + (T_2^{potRange} - T_2^{remBat})]$ and corresponding battery level of L_i^2, L_i^1 can achieve the same or a higher battery level at the same or an earlier start of service. An illustrative example is given in Fig. 2 (right plot). It depicts the tradeoff functions $L^l, l = 1, \ldots, 5$ of five labels. Function L^1 dominates function L^5, but does not dominate functions L^2, L^3, and L^4 because condition (9c) is violated in the case of L^2, condition (9d) in the case of L^3, and (9e) in the case of L^4.

4.1.2. Labeling algorithms for EVRPD-C variants

In this section, we detail the labeling algorithms for the variants with the objective of minimizing completion times and the three considered recharging strategies. For the EVRPD-C, we define $\tilde{c}_{ij}^k = -1/2\tilde{\pi}_i - 1/2\tilde{\pi}_j$ with the same definitions for $\tilde{\pi}_n, \tilde{\pi}_r, \tilde{\pi}_{o_k}$, and $\tilde{\pi}_{d_k}$ as in Section 4.1.1. Note the missing routing distance term c_{ij} in the definition of $\tilde{c}_{ij}^k$ compared to the EVRPD-D case. With these definitions, the reduced cost $\tilde{c}_p$ of a route $p \in \Omega_k$ is $\tilde{c}_p = \sum_{n \in N(p)} T_n + \sum_{(i,j) \in A(p)} \tilde{c}_{ij}^k$, where $N(p)$ denotes the customers visited on route p and T_n denotes the earliest possible start of service at a customer $n \in N(p)$.

The labeling algorithms for EVRPD-C and the different recharging strategies are all related to those for the corresponding EVRPD-D variant. While in the FR and BS cases only the REF for the reduced-cost resource has to be redefined, the required modifications in the PR case are much more involved.

EVRPD-C with full recharges. For the EVRPD-C(FR), all components of the labeling algorithm, i.e., label definition and resources, initial label, feasibility conditions (3), REFs (4) and dominance rule (5), are identical to the one for the EVRPD-D, except for the REF (4a) of the reduced-cost resource which is substituted by the REF

$$T_j^{cost} = \begin{cases} T_i^{cost} + T_j^{tMin} + \tilde{c}_{ij}^k & \text{if } j \in N, \\ T_i^{cost} + \tilde{c}_{ij}^k & \text{otherwise.} \end{cases} \tag{10}$$

EVRPD-C with battery swapping. Similar to FR, the labeling algorithms for the BS variants when minimizing routing distances and when minimizing the sum of completion times are almost identical. Again, we just need to substitute the reduced-cost REF (4a) by (10) in the latter. All other components remain.

EVRPD-C with partial recharges. We now describe the labeling algorithm for the EVRPD-C(PR). From an algorithmic point of view, this is clearly the most challenging variant considered. As with PR in the distance-minimizing variant, the amount to recharge at a recharging station has to be decided. Again, this can only be done a posteriori because it depends on the part of the route that succeeds the visit to the recharging station which is unknown for all labels not ending at the destination. In addition to the tradeoff between battery level and time, there is another tradeoff between the battery level and the cost when minimizing the sum of completion times: When recharging longer, all customers that are visited after the respective recharging station are serviced later, hence their service start time (and completion time) increases. Dealing with this additional tradeoff requires an additional resource compared to the distance-minimization variant. This resource is needed to model and propagate the new tradeoff function. Moreover, the definition of the reduced-cost resource, the reduced-cost REF, and the dominance rule have to be adapted to account for the more complex situation.

A label L_i involves the following resources $L_i = (T_i^{cost}, T_i^{load}, T_i^{tMin}, (T_i^{unreach_n})_{n\in N}, T_i^{nRch}, T_i^{remBat}, T_i^{potRange}, T_i^{nCust})$. Compared to the distance-minimization variant with PR, the additional resource T_i^{nCust} is used, and the definition of the reduced-cost resource T_i^{cost} is slightly changed. Their meaning is as follows:

T_i^{cost} the minimum reduced cost of path p. If a recharging station is visited on p prior to i, T_i^{cost} incorporates a minimum recharge, if necessary, to reach vertex i battery feasible. T_i^{cost} corresponds to the reduced cost of p when the start of service at vertex i is T_i^{tMin} and the battery level reaching vertex i is T_i^{remBat}.

T_i^{nCust} the number of customers visited on p between the last recharging station and vertex i (inclusive).

As in the distance-minimization case, the tradeoff between battery level and time is modeled with resources T_i^{tMin}, T_i^{remBat}, and $T_i^{potRange}$ (see Section 4.1.1). The tradeoff between battery level and reduced cost is modeled using resources T_i^{cost}, T_i^{remBat}, $T_i^{potRange}$, and T_i^{nCust}. It is illustrated in Fig. 3 (left plot). Note first that instead of the battery level, we could also use the time on the x-axis describing the tradeoff (recall that there is a one-to-one relation between the battery level and the time as given by the first tradeoff function). As by definition of the resources, the minimum reduced cost T_i^{cost} is obtained with the minimum battery level T_i^{remBat} or equivalently the earliest start of service T_i^{tMin}. To realize a higher battery level, the recharging duration at the preceding recharging station has to be increased. This implies that the visits at all customers after the recharging station are delayed by exactly the same amount of time. Hence, their service start times and the reduced-cost increase correspondingly. The slope of the tradeoff function is therefore equal to the number T_i^{nCust} of customers between the last recharging station and vertex i (inclusive).

The components of the initial label L_{o_k} at the origin o_k are initialized to the same values as in the distance-minimization variant, and the new resource T_i^{nCust} is initialized to zero. The feasibility conditions for label extensions and the REFs also remain, with the exception of the reduced-cost REF. Its updated version and the REF for T_i^{nCust} are given by:

$$T_j^{cost} = \begin{cases} T_i^{cost} + T_j^{tMin} - \tilde{c}_{ij}^k & \text{if } T_i^{nRch} = 0 \wedge j \in N, \\ T_i^{cost} - \tilde{c}_{ij}^k & \text{if } T_i^{nRch} = 0 \wedge j \notin N, \\ T_i^{cost} + T_j^{tMin} + \\ (T_j^{nCust} - 1) & \\ \quad \cdot \max\{0, h_{ij} - T_i^{remBat}\} - \tilde{c}_{ij}^k & \text{if } T_i^{nRch} > 0 \wedge j \in N, \\ T_i^{cost} + T_j^{nCust} & \\ \quad \cdot \max\{0, h_{ij} - T_i^{remBat}\} - \tilde{c}_{ij}^k & \text{if } T_i^{nRch} > 0 \wedge j \notin N, \end{cases} \tag{11a}$$

$$T_j^{nCust} = \begin{cases} 0 & \text{if } i \in R \wedge j \notin N, \\ 1 & \text{if } i \in R \wedge j \in N, \\ T_i^{nCust} + 1 & \text{if } i \notin R \wedge j \in N \wedge T_i^{nRch} > 0, \\ T_i^{nCust} & \text{otherwise.} \end{cases} \tag{11b}$$

Let $L^k = (T_k^{cost}, T_k^{load}, T_k^{tMin}, (T_k^{unreach_n})_{n\in N}, T_k^{nRch}, T_k^{remBat}, T_k^{potRange}, T_k^{nCust})$, $k \in \{1,2\}$, be two labels associated with different paths ending at the same vertex. Label L^1 dominates label L^2 if

$$T_1^{load} \le T_2^{load}, \tag{12a}$$

$$T_1^{tMin} \le T_2^{tMin}, \tag{12b}$$

$$T_1^{potRange} \ge T_2^{potRange}, \tag{12c}$$

$$T_1^{remBat} + T_2^{tMin} - T_1^{tMin} \ge T_2^{remBat}, \tag{12d}$$

$$T_1^{cost} + \max\{0, T_2^{remBat} - T_1^{remBat}\} \cdot T_1^{nCust} \le T_2^{cost}, \tag{12e}$$

$$T_1^{cost} + \max\{0, T_2^{potRange} - T_1^{remBat}\} \cdot T_1^{nCust} \le T_2^{cost} + (T_2^{potRange} - T_2^{remBat}) \cdot T_2^{nCust}, \tag{12f}$$

$$T_1^{unreach_n} \le T_2^{unreach_n} \quad n \in N. \tag{12g}$$

Again, conditions (12a) and (12g) are straightforward due to non-decreasing REFs. Conditions (12b)–(12d) ensure the dominance of L_i^1's first tradeoff function (battery level vs. time) over L_i^2's analog to the distance-minimization case. The dominance of L_i^1's second tradeoff function (battery level vs. reduced cost) over L_i^2's is guaranteed by conditions (12c), (12e), and (12f) in a similar fashion. For each potential battery level $H \in [T_2^{remBat}, T_2^{potRange}]$ and associated reduced cost of L_i^2, label L_i^1 can realize the same or a higher battery level with reduced cost that are not larger. The situation is exemplified in Fig. 3 (right plot), which shows the tradeoff functions $L^l, l = 1, \dots, 5$ of five labels. Function L^1 does not dominate functions L^2, L^3, and L^4 because conditions (12c), (12e), and (12f) are not fulfilled. Function L^1 does dominate function L^5.

4.1.3. Acceleration techniques

Since the solution of the pricing problem consumes by far the most computation time in the BPC, we use additional techniques to speed up the pricing process. First, we relax the elementary conditions of the routes and use the well-known *ng*-path relaxation (Baldacci, Mingozzi, & Roberti, 2011) resulting in pseudo-polynomially solvable pricing problems instead of the strongly NP-hard elementary case. For details on the general impacts on the BPC algorithm and on the necessary modifications of the labeling algorithms, we refer to Baldacci et al. (2011).

Second, partial pricing (Desaulniers, Desrosiers, & Solomon, 2002) is used to quickly identify negative reduced-cost routes. It is realized by subsequently performing pricing on a series of arc-reduced networks specified by a parameter κ as follows. Given κ, for each vehicle type $k \in K$ and vertex $i \in N \cup R$, the number of ingoing arcs $(i,j) \in A_k$ with $i \ne o_k$ is at least $\min\{\kappa, d_i^-\}$, where $d_i^- = |\{(j,i) \in A_k\}|$ is the in-degree of i in the full network. We always choose the arcs with the best reduced costs. The number of outgoing arcs is limited in the analog fashion. In addition, we keep all depot arcs $\{(i,j) \in A_k : i = o_k \vee j = d_k\}$. In each CG iteration, we execute the labeling algorithm for all vehicle types $k \in K$ and the same value of κ one after another. We start with the smallest value of κ and increase κ when no route with negative reduced cost is found for any of the vehicle types. As soon as one or more negative reduced-cost routes are found by the labeling algorithm for any combination of vehicle type and κ, they are returned to the RMP, and pricing is terminated for this iteration. In our experiments, we use the sizes $\kappa \in \{2, 5, 10, |V_k|\}$.

4.2. Cutting

To strengthen the linear relaxations of the RMP, several families of valid inequalities have been successfully applied in BPC approaches for VRPs. In our BPC for the EVRPD variants, we only use limited memory subset-row inequalities. In pretests, we also experimented with the well-known rounded capacity cuts (Naddef & Rinaldi, 2002) and 2-path cuts (Kohl, Desrosiers, Madsen, Solomon, & Soumis, 1999). Both

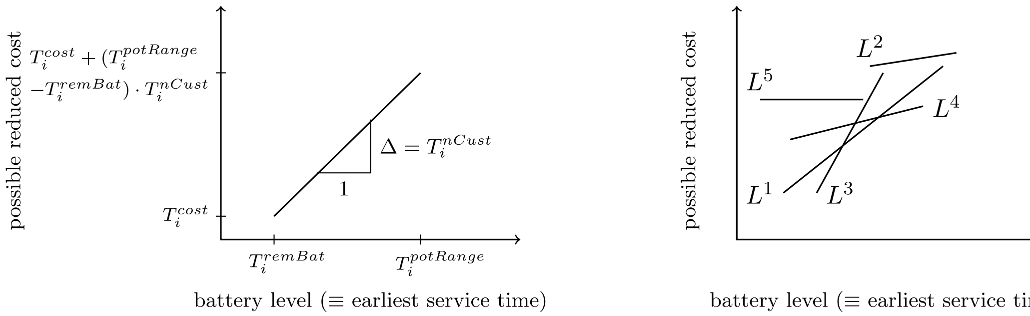

Fig. 3. Tradeoff between the battery level and the possible reduced cost.

families, however, did not prove helpful, so we do not use them in our BPC.

Subset-row inequalities (SRIs) were originally introduced by Jepsen, Petersen, Spoorendonk, and Pisinger (2008) for the VRP with time windows. Each SRI is defined on a subset $U \subset N$ of customers. As commonly done, we restrict ourselves to subsets of cardinality three. In this case, a subset-row inequality is given by $\sum_{k\in K}\sum_{p\in\Omega_k}\lfloor\frac{u_p}{2}\rfloor\theta_p \leq 1$, where u_p denotes the number of times route p visits customers in U. Violated SRIs can be separated by straightforward enumeration.

SRIs are non-robust cuts. Each SRI requires an additional binary resource in the labeling algorithms for solving the pricing problems, thus making pricing harder. To reduce the negative impact on the pricing problems, Pecin, Pessoa, Poggi and Uchoa (2017) have proposed a limited memory variant of the SRIs. Their idea is similar to the *ng*-route relaxation: A vertex memory controls at which vertices the state of an SRI resource remains relevant. We use the same separation algorithm and vertex memory as described in Pecin, Contardo, Desaulniers and Uchoa (2017).

4.3. Branching

Let $(\bar{\theta}_p)_{p\in\Omega_k,k\in K}$ be a fractional solution of the RMP (1) and denote by $\bar{b}_{rp}$ and $\bar{b}_{ijp}$ the numbers of times a route $p \in \Omega_k, k \in K$ visits recharging station $r \in R$ and uses arc $(i,j) \in A$, respectively, in this solution. To ensure integer solutions, we apply the same hierarchical branching scheme as in Desaulniers et al. (2016) slightly adapted to the heterogeneous case.

First, we branch on the number of routes $\sum_{p\in\Omega_k}\bar{\theta}_p$ of a vehicle type $k \in K$. Second, we branch on the total number of recharges $\sum_{r\in R}\sum_{p\in\Omega_k}\bar{b}_{rp}$ of a vehicle type $k \in K$. Third, we branch on the number of recharges $\sum_{p\in\Omega_k}\bar{b}_{rp}$ of a vehicle type $k \in K$ at a recharging station $r \in R$. Finally, we branch on the total flow $\sum_{k\in K}\sum_{p\in\Omega_k}\bar{b}_{ijp}$ on an arc $(i,j) \in A$.

For all levels, we always select a branching variable for which the fractional value is closest to 0.5. We then create two branches, one forcing the branching variable to be larger or equal to the rounded-up value and the second one forcing it to be smaller or equal to the rounded-down value. Branching decisions of the first three levels are imposed by adding a linear inequality to the RMP, while those of the fourth level are imposed by temporarily removing infeasible arcs from the pricing problem networks. Moreover, in all cases, all routes in the RMP that are not compatible with the branching decisions are temporarily forced to zero. Note that all branching decisions leave the pricing problems structurally unchanged.

We explore the search tree using a best-bound-first strategy for the selection of the next branch-and-bound node.

5. Computational study

The aim of our computational study is twofold. First, we analyze the computational performance of our BPC. Second, we generate managerial insights on the AMR route planning problem with respect to the different recharging strategies and compare the different objective functions. To this end, we consider two different benchmark sets: Our first benchmark set (`Sol`) is derived from a standard benchmark from the literature for road transportation with electric trucks (Schneider et al., 2014). This set is primarily used to assess the computational performance of our algorithm on well-established instances. Our second set (`Karis`) introduces new instances originating from a real-world pilot study in which AMRs were used to supply an assembly line for sports cars. These instances feature typical characteristics of AMRs in intralogistics, which differ considerably from road transportation with electric trucks. For example, in the former, we often see:

- short driving times due to short distances,
- short recharging durations due to small battery capacities,
- initial battery level might be low (in contrast, trucks are usually recharged over night),
- a layout with optimized positions for recharging stations on links of important sink–source relations,
- a rather regular network, usually without differentiation of street types and related speeds,
- starting/ending locations of vehicles are distributed,
- symmetric transport jobs with exchange of full and empty trolleys.

We published all instances and their corresponding solutions on a Zenodo repository (Meyer et al., 2024).

The BPC algorithm was implemented in C++ and compiled into 64-bit single-thread code with MS Visual Studio 2015. CPLEX 12.9 with default parameters (except for the time limit and allowing only a single thread) was used to reoptimize the RMPs. The computations were carried out on a standard PC running Windows 10 with an Intel(R) Core(TM) i7-6900k processor clocked at 3.2 GHz and equipped with 64 GB RAM main memory. The computational time limit was set to 3600 s.

5.1. Benchmark instances

`Sol` instances. The `Sol` benchmark is based on the EVRPTW instances introduced in Desaulniers et al. (2016) and Schneider et al. (2014) that extend the well-known Solomon benchmark instances for the VRPTW (Solomon, 1987). They reflect a pure distribution or pick-up use case with electric trucks on a street network. For the problem variants addressed, we adapt the instances as follows: First, we ignore the lower time window limits as we only consider due dates and, therefore, only need the upper time window limits. Second, we set a reasonable time for BS. Third, we add reasonable bounds on the fleet size. In the following, we briefly sketch the main characteristics of the instances and detail our adaptations.

The original Solomon benchmark comprises instances with different geographical distribution of the customers (customers clustered across the area of consideration, randomly distributed customers, a mixture of both) and different lengths of the scheduling horizon (short horizon with narrow time windows, long horizon with wide time windows).

Table 2
The number of instances solved to optimality in pre-tests.

#Cust	BS	FR	PR
25	156	165	165
50	117	120	126
All	273	285	291

In all instances, an unlimited fleet of homogeneous vehicles is available at a common depot. The 100-customer Solomon instances were adapted to the EVRPTW by Schneider et al. (2014). For details on how the recharging stations are placed and how the battery capacity, the consumption rate, the recharging rate, and the time windows are set in the adapted instances, we refer to Schneider et al. (2014). Please note that the battery-related parameters were created for algorithmic benchmarking at the instance level. They do not reflect the technological characteristics of electric trucks. For example, the charging rate was set in relation to customer service times. The homogeneous vehicles all start with a full battery and are allowed to return to the depot with an empty battery. For our study, we use the mid-sized instances introduced by Desaulniers et al. (2016) containing only the first 25 or 50 customers of each instance.

Typically, BS is considerably faster than recharging. For our instances, we set the duration b for BS to 10% of the maximum battery capacity H_k^{max}, i.e., the duration for recharging the battery from 0% to 100% load, to work with a scalable value.

If completion times are minimized, the fleet size needs to be limited. Otherwise, a solution in which each customer is serviced as early as possible by a dedicated vehicle is trivially optimal. To generate reasonable bounds on the fleet size for each instance, we solved the EVRPD-D with the three recharging strategies and an unlimited fleet in a preprocessing step and imposed the number of resulting vehicles as the fleet size limit (v0 instances). To vary the number of vehicles, we also consider instances with one (v1) or two (v2) additional available vehicles. If no optimal solution could be found in the preprocessing phase, the instance was not considered any further. Table 2 shows the number of resulting instances per recharging strategy and instance size given by the number of customers (#cust).

To interpret the results of the computational study correctly, it is important to understand that our fleet limit procedure may result in a different number of vehicles for the same instance but different recharging strategies. This design assures that recharging decisions play an important role in all recharging setups: The number of necessary vehicles in case of FR is the highest. If this number is taken as fleet size limit for the other two strategies, preliminary tests showed that most tours for BS and PR only contained one recharging station, and studies on the effects of recharging decisions would not be meaningful. As a result, we can only compare the solutions for the different objectives but not for the different recharging strategies. To compare the effects of different recharging strategies, we designed the `Karis` instances introduced next.

`Karis` instances. This set of instances originates from the KarisPro Audi Sport pilot study, in which the AMR KarisPro (see Fig. 4) was tested in a production site of the Audi Sport GmbH (KARIS PRO Consortium, 2016). In this setup, the sports car Audi R8 was produced in a synchronized assembly with 16 cycles and a cycle time of approximately 20 min. KarisPro AMRs took over the sequenced delivery of vehicle-specific material trolleys between the picking area and the assembly line, i.e., they transported full trolleys from the picking area to the assembly line and returned empty trolleys. Before, employees took over the transport of 33 of these trolleys in each cycle. The pilot study was conducted in 2016 with the research prototypes of the AMR KarisPro. The AMR is now commercially available (GEBHARDT Intralogistics Group, 2022).

Commonly in intralogistics, AMRs have no fixed depot. They either wait at different recharging stations or at positions where pickup jobs arrive. In our study, the AMRs start and return to different pickup stations of the picking area. Hence, we have to handle a heterogeneous fleet with different origins and destinations.

Table 3
Factor variations of `Karis` instance set.

Factor	Variants
Scenario size (per cycle)	S: 6 assembly stations, 12 jobs, 4 vehicles M: 12 assembly stations, 24 jobs, 8 vehicles L: 18 assembly stations, 36 jobs, 12 vehicles XL: 24 assembly stations, 48 jobs, 16 vehicles
Number of cycles	1, 2, 3, 4
Cycle time	30 min
Battery capacity	39 min (duration for recharging battery from 0% to 100% load)
Battery swapping duration	5 min
Recharging speed	FAST: 0.25 (20 s for 1 min driving) SLOW: 0.5 (30 s for 1 min driving)
Recharging state	EMPTY: initial and final level 20% FULL: initial and final level 80% SHIFT: initial level 80%, final level 20%
Battery capacity variants	25%, 50%, 75%, 100% (of original battery capacity)
Recharging station	One recharging station between picking area and assembly line with a capacity corresponding to the number of vehicles

To cover a broad context of industrial applications and instance sizes, we varied the setup of the pilot application and the vehicle characteristics. Table 3 gives an overview of the factors determining the instances and the variants we derived.

The scenario size determines the number of assembly stations, the number of jobs per station and cycle, and the number of available vehicles. The jobs either correspond to the delivery of a full trolley or the pickup of an empty trolley. The scenario size S corresponds to the pilot study with six assembly stations, twelve jobs, and four vehicles. We added three more variants comprising up to 48 jobs and up to 16 vehicles referred to as M, L, and XL (see Table 3). We furthermore varied the number of cycles from one to four, assuming a fixed cycle time of 30 min. If we consider the S scenario size and two cycles, four vehicles have to cover 24 jobs (12 jobs in each cycle) within 60 min (twice the cycle time of 30 min).

The battery-related parameters are set as follows: We assume a battery capacity H_k^{max} of 39 min (duration to fully recharge an empty battery) as a base value. Additional instances are considered by reducing H_k^{max} to 25%, 50%, and 75% of the original value of 39 min. For BS, we assume a battery swapping duration b of 5 min. We consider two recharging speeds: FAST and SLOW. In the former, for one minute of driving, 20 s of recharging is necessary. In the latter, the corresponding recharging duration is 30 s. SLOW was valid for the KarisPro research prototype. However, in the meantime, the recharging technologies have improved considerably. We can calculate h_{ij} by multiplying t_{ij} with the respective recharging speed. The recharging state describes the initial battery level (H_k^{init}) and the final battery threshold (H_k^{end}) of the vehicles as percentage of the battery capacity. In the EMPTY (FULL) instances, the initial battery level of the vehicles is 20% (80%), and the final battery level must not be below 20% (80%). For SHIFT, the vehicles start with almost full batteries (80%) and are allowed to finish almost empty (20%). This scenario assumes that the vehicles can be recharged, e.g., overnight.

As in the pilot case study, we assume one recharging station located between the picking area and the assembly line and whose capacity corresponds to the number of available vehicles. Distances, driving times, and energy consumption have been derived from the layout and the vehicle characteristics of the KarisPro research prototype AMR.

In our computational study, we limit the scenario size and number of cycle combinations to the following: For battery capacity of 100%, $\{S, M, L\} \times \{1, 2, 3, 4\}$ and $\{XL\} \times \{1, 2\}$, for battery capacities of 25%,

Fig. 4. The KarisPro AMR research prototype delivers a material transport trolley from the picking area to the assembly line.

Table 4
Overview of Karis instance set.

Scenario	Cycles	#cust	#veh	#inst by battery capacity variants			
				25%	50%	75%	100%
S	1	12	4	6	6	6	6
	2	24	4	6	6	6	6
	3	36	4	6	6	6	6
	4	48	4	6	6	6	6
M	1	24	8	6	6	6	6
	2	48	8	6	6	6	6
	3	72	8	–	–	–	6
	4	96	8	–	–	–	6
L	1	36	12	6	6	6	6
	2	72	12	6	6	6	6
	3	108	12	–	–	–	6
	4	144	12	–	–	–	6
XL	1	48	16	6	6	6	6
	2	96	16	6	6	6	6

50%, and 75%, $\{S\} \times \{1,2,3,4\}$ and $\{M, L, XL\} \times \{1,2\}$. All other factors are treated according to a full factorial design. In total, this results in 264 instances that we solved for all objective-recharging strategy combinations. This results in 264 instances that we solved for all objective-recharging strategy combinations.

Table 4 provides an overview of the Karis instance set. In particular, it highlights the main information regarding instance size, i.e., number of jobs transformed to customers (#cust) and number of vehicles (#veh), of the corresponding transformed EVRPD. A complete overview of instance characteristics is given in Appendix B.

5.2. Results for the Sol instances

In this section, we briefly discuss the algorithmic performance of our BPC variants on the Sol benchmark set and focus on the comparison of the results considering the two different objective variants. Please remember that the results of different recharging strategies are not comparable as the number of vehicles for the same instances varies (see Section 5.1).

5.2.1. Algorithmic performance

Table 5 summarizes the algorithmic performance of our BPC variants. It provides the number of instances (#inst), the percentage of instances solved to proven optimality within the time limit (opt [%]), the total run time (rt [s]), and the run time for solving the root node (rtr [s]) in seconds for the instances solved to optimality, and the root node gap (gap [%]) and number of branch-and-bound nodes solved (#nds) for the instances solved to optimality. Consistently over all recharging strategies, the BPC is able to solve to optimality slightly more instances for the DI objective (76%–79%) compared to the CT objective (66%–69%). This can be attributed to the harder-to-solve pricing problems in the CT case, which can be seen from the overall much longer computation times for the root node. This effect is especially pronounced for PR, where for CT only a much weaker dominance rule than for DI is applicable in the labeling algorithm. Root node gaps are very small in general and consistently smaller for the CT objective, leading to smaller branch-and-bound trees compared to DI.

5.2.2. Comparing objectives

To analyze the impact of the different objective functions, we filter the results to the 133 instances for which all six variants – two objectives with three recharging strategies – are solved to optimality. Clearly, DI objective solutions result in shorter total distances but higher total completion times compared to CT, and vice versa. The percentage deviations are shown in Table 6. Over all recharging strategies, the total traveled distance increases on average by around 23% with a maximum of 60% if objective CT is applied, while the average increase in completion times when using the DI objective compared to using the CT objective is as high as 50% to 60% with a maximum of more than 200%.

We can also observe structural differences in the resulting solutions of the different objective functions and recharging strategies. For the sake of brevity, we only summarize the main characteristics in the following. A more detailed analysis is provided in Appendix C.

Used vehicles: CT solutions always schedule all vehicles, as CT incentivizes short routes with early completion times and uses the initial energy level of the whole fleet. In DI solutions, the extra vehicles of v1 and v2 are not used to avoid the extra distance from and to the common depot of the vehicles and create longer routes to increase the probability of good customer combinations.

Number of recharging stops: CT solutions tend to have slightly more recharging stops than DI solutions.

Timing of recharging stops: In CT solutions, recharging stops are scheduled later than in DI solutions. A large share of routes contains a recharging stop just before returning to the depot.

Duration of recharging: The recharging duration mainly depends on the recharging strategy. The insights are similar to the Karis benchmark and are discussed in the following section.

Table 5
Performance overview of the BPC on the Sol benchmark.

#cust	Strategy	#inst	DI objective					CT objective				
			opt[%]	rt[s]	rtr[s]	gap[%]	#nds	opt[%]	rt[s]	rtr[s]	gap[%]	#nds
25	BS	156	86	123.4	43.4	0.012	6.8	79	134.1	100.4	0.004	3.3
	FR	165	81	183.9	87.2	0.011	5.4	79	161.4	91.0	0.006	11.3
	PR	165	80	199.0	13.8	0.013	8.0	84	197.3	159.0	0.004	5.0
50	BS	117	69	809.7	214.9	0.018	26.7	50	659.4	316.8	0.005	8.5
	FR	120	70	607.5	227.2	0.018	26.9	56	520.4	224.2	0.009	19.1
	PR	126	71	478.1	166.1	0.017	24.3	48	654.3	237.9	0.007	15.0
All	BS	273	79	381.9	108.0	0.014	14.3	66	302.4	169.8	0.005	5.0
	FR	285	76	347.1	141.2	0.014	13.7	69	282.9	136.1	0.007	13.9
	PR	291	76	311.4	75.1	0.014	14.6	68	335.8	182.9	0.005	8.0

Table 6
Differences in total distance and completion time on the Sol benchmark.

	DI objective			CT objective		
	BS	FR	PR	BS	FR	PR
Mean total distance				+23.4%	+22.6%	+24.0%
Min total distance				+1.1%	+4.4%	+2.8%
Max total distance				+58.9%	+51.7%	+59.8%
Mean total completion time	+56.1%	+66.0%	+62.6%			
Min total completion time	+1.7%	+8.9%	+3.3%			
Max total completion time	+235.7%	+241.4%	+235.6%			

Table 7
Performance overview of the BPC on the Karis benchmark.

Objective	#inst	opt[%]	inf[%]	found[%]	unkn[%]	rt[s]	rtr[s]	gap[%]	#nds
DI	792	43.2	20.3	5.4	31.1	254.6	4.1	0.003	126.2
CT	792	54.4	20.3	8.2	17.0	179.8	18.1	0.001	30.8

5.3. Results for the Karis instances

In this section, we detail our computational results for the Karis benchmark, focusing on an in-depth comparison of the obtained solutions to generate managerial insights on the AMR route planning problem with respect to the different recharging strategies and objective functions.

5.3.1. Algorithmic performance

Table 7 summarizes our results aggregated by objective function. The additional columns have the following meaning: the percentage of provably infeasible instances (inf[%]), i.e., instances for which we could prove that no feasible solution exists, the percentage of instances for which a feasible solution without proven optimality could be found (found[%]), and the percentage of instances for which no feasible solution has been found within the time limit (unkn[%]). Overall, and in contrast to the Sol instances, we can see that the CT variant seems to be easier compared to DI: More instances are solved to proven optimality, and the average run time is smaller. This behavior can mainly be attributed to the smaller gaps leading to much smaller search trees in the CT case. The times for solving the root node, on the other hand, are much larger for CT, indicating that the corresponding pricing problems are harder to solve compared to those of the DI objective.

For both objectives, our BPC was able to determine feasibility (*opt* or *found*) or infeasibility for the majority of instances (83% for CT, 69% for DI). Fig. 5 shows the percentage of instances with known status (feasible or infeasible) of the two objectives over the number of customers. Our BPC reaches a high share of instances with known status for up to 48 customers per instance, especially in the case of CT and irrespective of the recharging strategy. However, the numbers drop considerably below 50% if we include 96 and more customers. Another factor influencing the hardness of the instances can be read from Table 8. It provides an overview of solutions for instances with 48 customers resulting from different combinations of the number of cycles (Cyc) and the scenario sizes (Scenario). Apparently, the instances become harder when the number of cycles and, hence, the planning horizon increases leading to generally longer routes. This effect is not surprising, but it is much more pronounced for the DI objective. This can be attributed to the fact that in CT always all vehicles are used to obtain the shortest possible routes with small completion times, while DI incentivizes long routes using a low number of vehicles to minimize distances.

Table 8
Overview of solution status for instances with 48 customers.

Objective	Cyc	Scenario	#inst	opt[%]	infeas[%]	found[%]	unkn[%]
DI	1	XL	72	55.6	44.4	0.0	0.0
	2	M	72	63.9	18.1	16.7	1.0
	4	S	72	2.8	6.9	1.4	89.0
CT	1	XL	72	48.6	44.4	6.9	0.0
	2	M	72	62.5	18.1	15.3	4.0
	4	S	72	70.8	6.9	5.6	17.0

5.3.2. Comparing objectives

For comparing the results of the different objectives, we filtered the instances to those for which we found an optimal solution for both objective variants. In the case of BS, 139 instances are left, in the case of FR 91, and in the case of PR 134. Table 9 shows the mean, minimum, and maximum increase of the total distance and the total completion time if the respective other objective is used. It can be seen that differences in this AMR scenario are smaller than for the Sol benchmark. The total distance traveled increases by about 10% to 15% on average (with maximum of around 30%) when using the CT objective, while the average increase in completion time when using the DI objective ranges from 33% for FR to around 50% for BS. Maximum values are consistently above 100% for the latter case.

5.3.3. Comparing recharging strategies

We now perform an in-depth comparison of the recharging strategies by looking at the feasibility of the instances, the recharging behavior, and the solution quality.

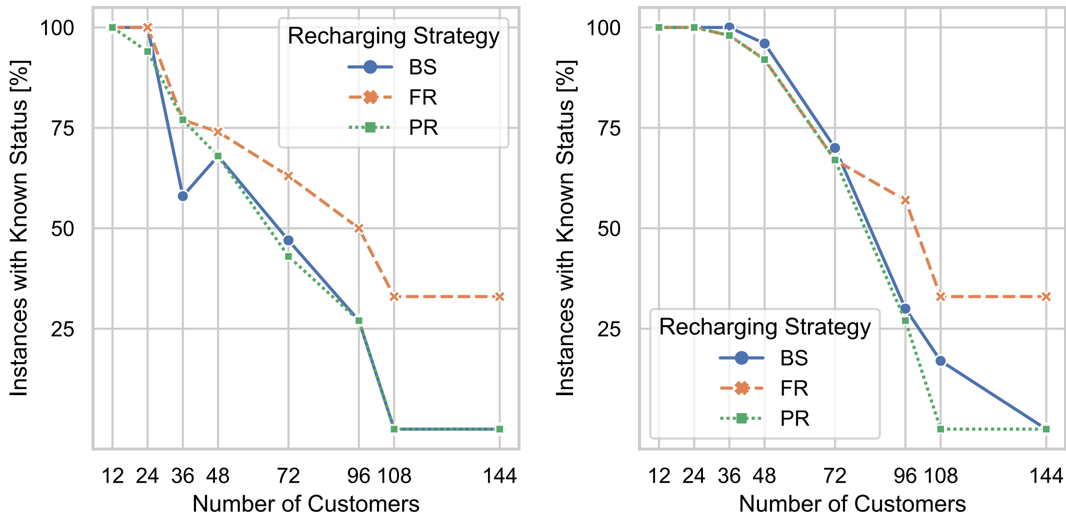

(a) DI objective.

(b) CT objective.

Fig. 5. Percentage of instances with known status over the number of customers.

Table 9
Differences in total distance and completion time on the Karis benchmark.

	DI objective			CT objective		
	BS	FR	PR	BS	FR	PR
Mean total distance				+14.0%	+11.7%	+13.5%
Min total distance				+1.8%	+0.1%	+0.7%
Max total distance				+29.6%	+28.3%	+26.7%
Mean total completion time	+52.3%	+33.5%	+43.0%			
Min total completion time	+6.7%	+3.7%	+5.9%			
Max total completion time	+148.7%	+176.7%	+183.9%			

Table 10
Overview of solution status by recharging strategy for CT variant for instances with 100% battery capacity.

Strategy	#inst	Opt	Found	Infeasible	Unknown
BS	84	51	7	0	26
FR	84	25	5	33	21
PR	84	42	6	7	29

Feasibility. Table 10 shows the solution status by recharging strategy for the considered 84 instances with a battery capacity of 100%. As the objective does not affect feasibility, we only show the results for CT. If BS is applied, no instance is provably infeasible. If IR is applied, on the other hand, several instances turn out to be infeasible. With an FR strategy, 33 out of the 84 instances (39%) are infeasible. If an PR strategy is applied, the number decreases to only seven instances (8%). This shows that in the case of IR, a larger AMR fleet size compared to BS is necessary to remain feasible. However, the PR strategy is able to considerably reduce this drawback of IR.

We next analyze the impact of the battery on feasibility, also considering the instances with a battery capacity of 25, 50, and 75% of the original capacity. In Fig. 6(a), we compare the number of infeasible instances out of the 41 instances for which there is no unknown solution status over all considered battery capacities and recharging strategies. Interestingly, not all recharging strategies benefit from a larger battery capacity. While for BS and PR the number of infeasible instances decreases with a larger battery, for the FR strategy, the number of infeasible instances stays at a very high level and even increases if the battery capacity goes up from 75 to 100%. This somewhat counter-intuitive behavior can be easily explained. An increased battery capacity results in a longer average recharging duration for FR, which in turn may imply that some deadlines can no longer be met. Fig. 6(b) highlights that this effect is particularly pronounced if the vehicles start with an empty battery (EMPTY) and have to recharge relatively early during the route. However, even if the vehicles start with a full battery (FULL), the increase in battery capacity from 75 to 100% results in more infeasible solutions.

Recharging Behavior. We analyze the recharging behavior of the different recharging strategies separately for each objective. Also, we filter the instances to the ones for which optimal solutions are known for all recharging strategies. In the CT case, this leaves us with 100 instances: 6 for EMPTY, 39 for FULL, and 55 for SHIFT.

Fig. 7(a) depicts the total number of recharges per instance differentiated by the recharging state. As expected, recharging plays a major role in scenarios EMPTY and FULL in which the vehicles start and may return with almost empty and full batteries, respectively. In these scenarios, the average number of recharges per route is above one, irrespective of the recharging strategy, see Fig. 7(b). For scenario SHIFT, where the vehicles start almost full and may return almost empty, recharging plays a minor role with an average number of recharges per route of less than 0.5.

An overview of the total time spent for recharging and the average recharging duration is provided in Figs. 7(c) and 7(d), respectively. Fig. 7(e) gives the total recharging duration relative to the total route duration. Note that for BS, the average recharging time equals the duration $b = 300$ s for swapping the battery. The total recharging duration makes the advantage of the BS technology obvious. Over all three scenarios, PR needs, on average, 2.2 times longer for recharging while FR even needs 3.7 times longer. This advantage does not only affect feasibility (see above) but also the total completion times (see below). The only exception is in scenario SHIFT, where the total and even the average recharging duration for PR is lower. Here, PR solutions use the flexibility to recharge more often but even shorter than 300 s on average. This picture is reconfirmed by Fig. 7(e).

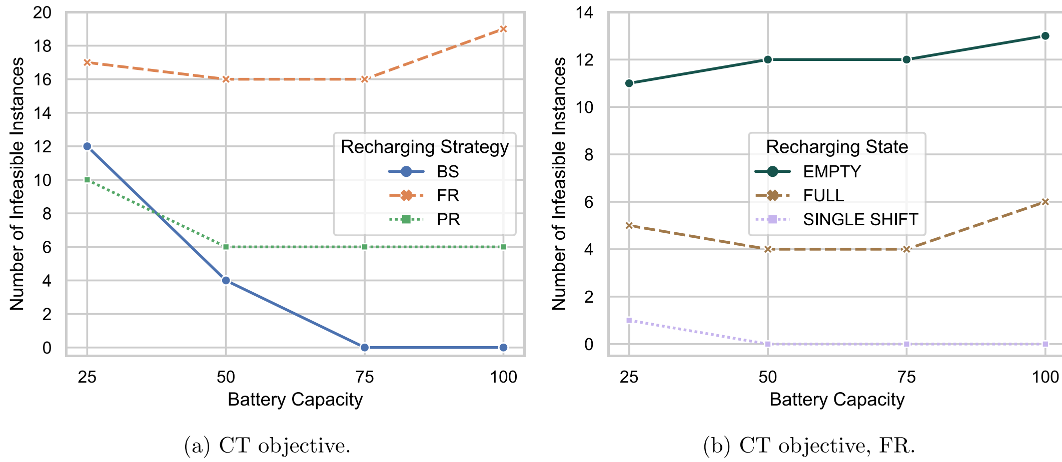

Fig. 6. Number of infeasible instances over the battery capacity by recharging strategy (a) and by recharging state in case of FR (b) for the 41 instances with known status for all considered variants.

The figures also clearly unveil the shortcomings of the FR strategy: As expected, the average recharging duration lies considerably above those of the other two strategies. Even more, also the total recharging duration is, on average, 1.7 times longer over all recharging states than for PR. This means that the battery level at the end of the planning horizon is considerably higher than necessary, especially in the EMPTY and SHIFT scenarios, in which vehicles are allowed to return almost empty.

We can also see that the additional flexibility of PR compared to FR is intensively used in the optimal solutions. This is especially true for the EMPTY scenario for which the number of recharges is highest, but the average recharging duration is comparatively short. For reaching a low total completion time, many but short recharging stops are scheduled.

For the DI objective, the number of instances used in this analysis, i.e., for which optimal solutions are known for all recharging strategies, are 3 for EMPTY, 19 for FULL, and 48 for SHIFT. Fig. 8 corresponds to Fig. 7 for the DI objective. Figs. 8(a) and 8(b) show that both the total number of recharges and the average number of recharging stops per route are lower than in the CT case. This confirms that DI incentivizes to avoid recharging stops in order to avoid detours to the recharging stations. In contrast, the effects considering the sum of recharging durations, the average recharging duration, and recharging duration relative to the route duration as depicted in Figs. 8(c), 8(d), and 8(e), respectively, are very similar to those observed for the CT objective.

In Fig. 9, we analyze the structure of the resulting routes of all optimal solutions. The graphs provide the following information. On the x-axis, we depict the positions of all recharging stops relative to the number of stops of the corresponding route. On the y-axis, the recharging durations of all recharging stops relative to the planning horizon of the corresponding instance are given. Recall that for BS a constant battery swap duration b is assumed. The differences in the chart result from different planning horizons and should not be interpreted otherwise.

The graphs illustrate that the objective mainly influences the distribution of the relative positions of the recharging stops. With a CT objective, clearly, left-skewed distributions are obtained compared to the considerably more symmetric distributions with DI. The relative recharging durations, on the other hand, are determined by the recharging strategy. With both objectives, PR allows considerably shorter recharging durations than its FR counterpart.

Solution quality. To compare solution quality, we look at the relative deviations of the resulting total distance and total completion time of FR and PR compared to BS. Table 11 provides an overview of these deviations when the CT and DI objectives are applied. To maintain comparability, we filtered the instances to the 100 (70) instances for which we found an optimal solution applying the CT (DI) objective for all three recharging strategies.

Table 11
Relative deviations of total distances and total completion times compared to BS solutions over 100 (70) instances for which an optimal solution was found for all three recharging strategies and CT (DI) objective.

Objective	Strategy	#inst	Relative deviation DI			Relative deviation CT		
			Mean	Min	Max	Mean	Min	Max
CT	FR	100	−0.3%	−8.6%	9.9%	8.4%	0.0%	134.6%
CT	PR	100	0.6%	−3.0%	11.7%	0.6%	−5.2%	17.6%
DI	FR	70	1.6%	0.0%	10.5%	1.0%	−35.3%	38.8%
DI	PR	70	0.8%	−0.8%	4.7%	−2.7%	−35.2%	15.4%

Table 11 indicates that the recharging strategy only has a rather small impact on the resulting distances if the DI objective is applied. For FR, the total distance is, on average, 1.6% higher with a maximum of 10.5%. For PR, the increase is 0.8% on average, with a maximum of 4.7%. In some cases, PR even enables shorter distances than BS.

The picture is different if we compare the total completion time applying the CT objective. In this case, FR solutions are, on average, more than 8% worse than BS solutions with a maximum of 134%. In contrast, the CT objective value of the PR strategy is, on average, only 0.6% higher than that of the BS strategy with a maximum of 18% and, hence, seems to be almost competitive. The more detailed analysis in Fig. 10, which provides the percentage deviation for the different recharging state scenarios, reveals that the impact of the recharging strategy on the total completion time is most pronounced for the EMPTY scenario, in which recharging is most relevant. There, the total completion time of FR is, on average, more than 50% higher compared to BS. When looking at total distances while applying the CT objective, the deviations of FR and PR compared to BS are with −0.3% and +0.6%, respectively, almost negligible.

6. Impact of recharging strategies

To the best of our knowledge, exact solutions for practical instances of AMR route planning in an industrial context have been generated for the first time in this paper. Our results allow the comparison of two different objective functions and, more importantly, different recharging strategies in a realistic manufacturing scenario. In the following, we summarize the main insights:

- BS decouples recharging and route execution and can generally be considered the most time-efficient recharging technology. This

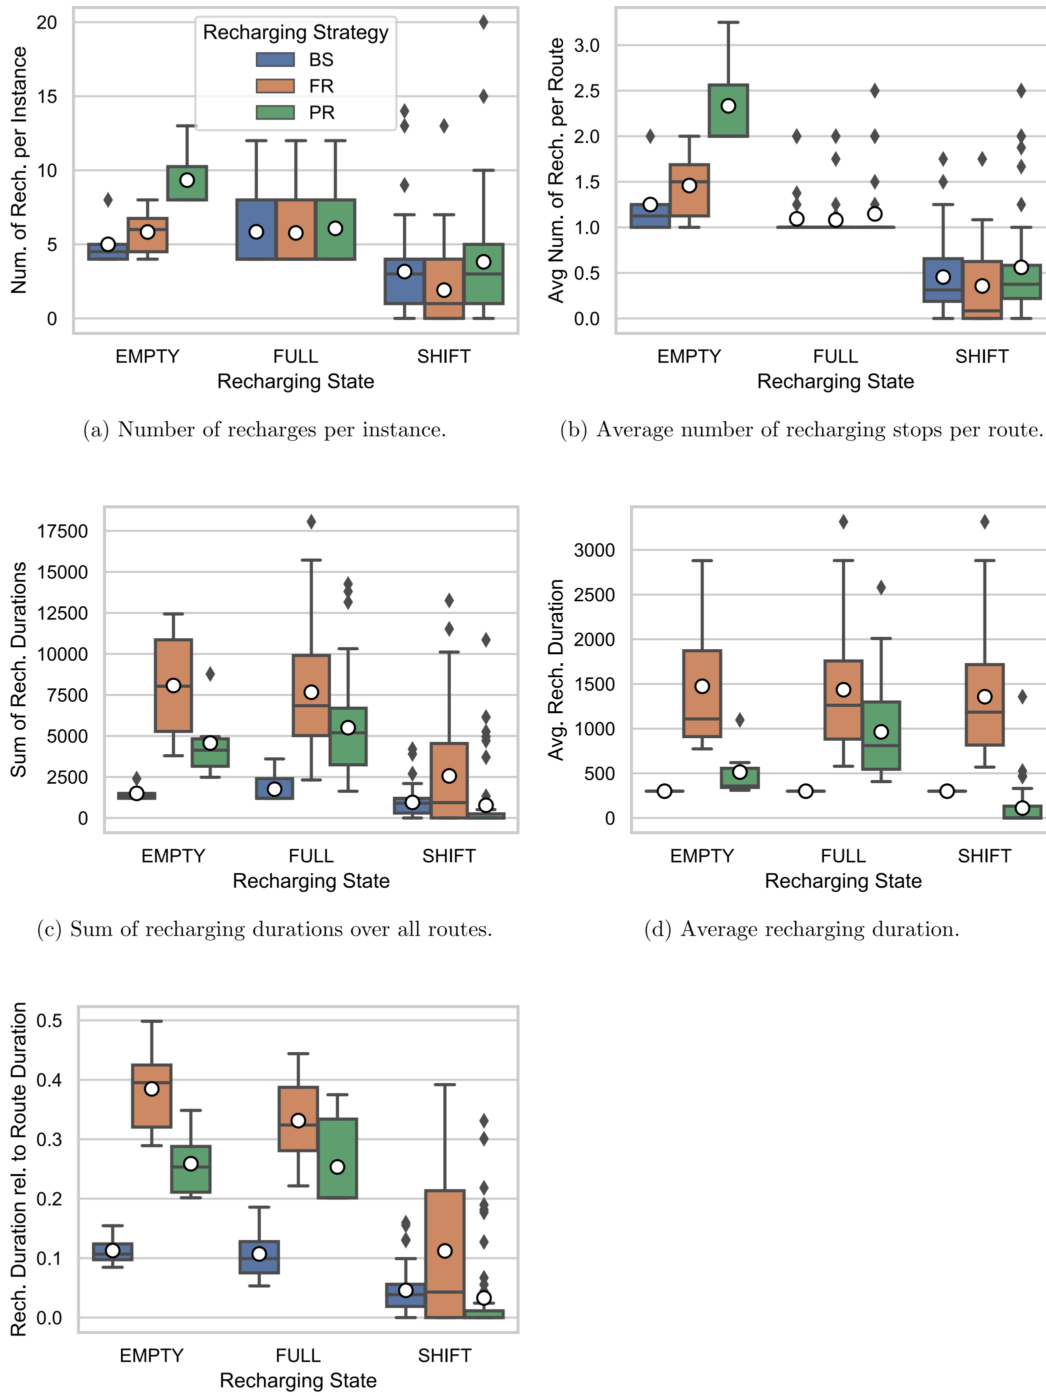

(a) Number of recharges per instance.

(b) Average number of recharging stops per route.

(c) Sum of recharging durations over all routes.

(d) Average recharging duration.

(e) Total recharging duration relative to route duration.

Fig. 7. Recharging behavior over 100 instances for which an optimal solution was found for the CT objective by recharging strategies.

is confirmed by our study. BS results in solutions with lower total completion time than IR in almost all of our scenarios. This comes at the cost of more expensive infrastructure and a larger number of batteries, which belong to the most expensive components of AMRs. Furthermore, the time-efficiency advantage of BS decreases if recharging in known off-peak hours, e.g., after a shift, is possible.

- An FR strategy, which is often proposed as a simplifying heuristic if IR is applied, performs poorly in our AMR scenario. A large number of infeasible instances can, in practice, only be countered with a larger vehicle fleet. Furthermore, in cases where recharging plays an important role, e.g., scenario EMPTY, solutions show a 50% higher average total completion time compared to BS. Even more and regardless of the results for our specific scenario, the effect that a larger battery capacity of the vehicles leads to more infeasibilities is not only counter-intuitive but, in our opinion, unacceptable.
- Surprisingly, for our AMR scenario, PR proved to be almost competitive to BS. Its solutions show only a few infeasibilities, and the average total completion time is, on average, below 4%, even in the EMPTY scenario. The results also highlight that the additional degree of freedom of deciding on the recharging duration compared to FR is extensively used. A route structure with many short recharging stops and a tendency to late recharges

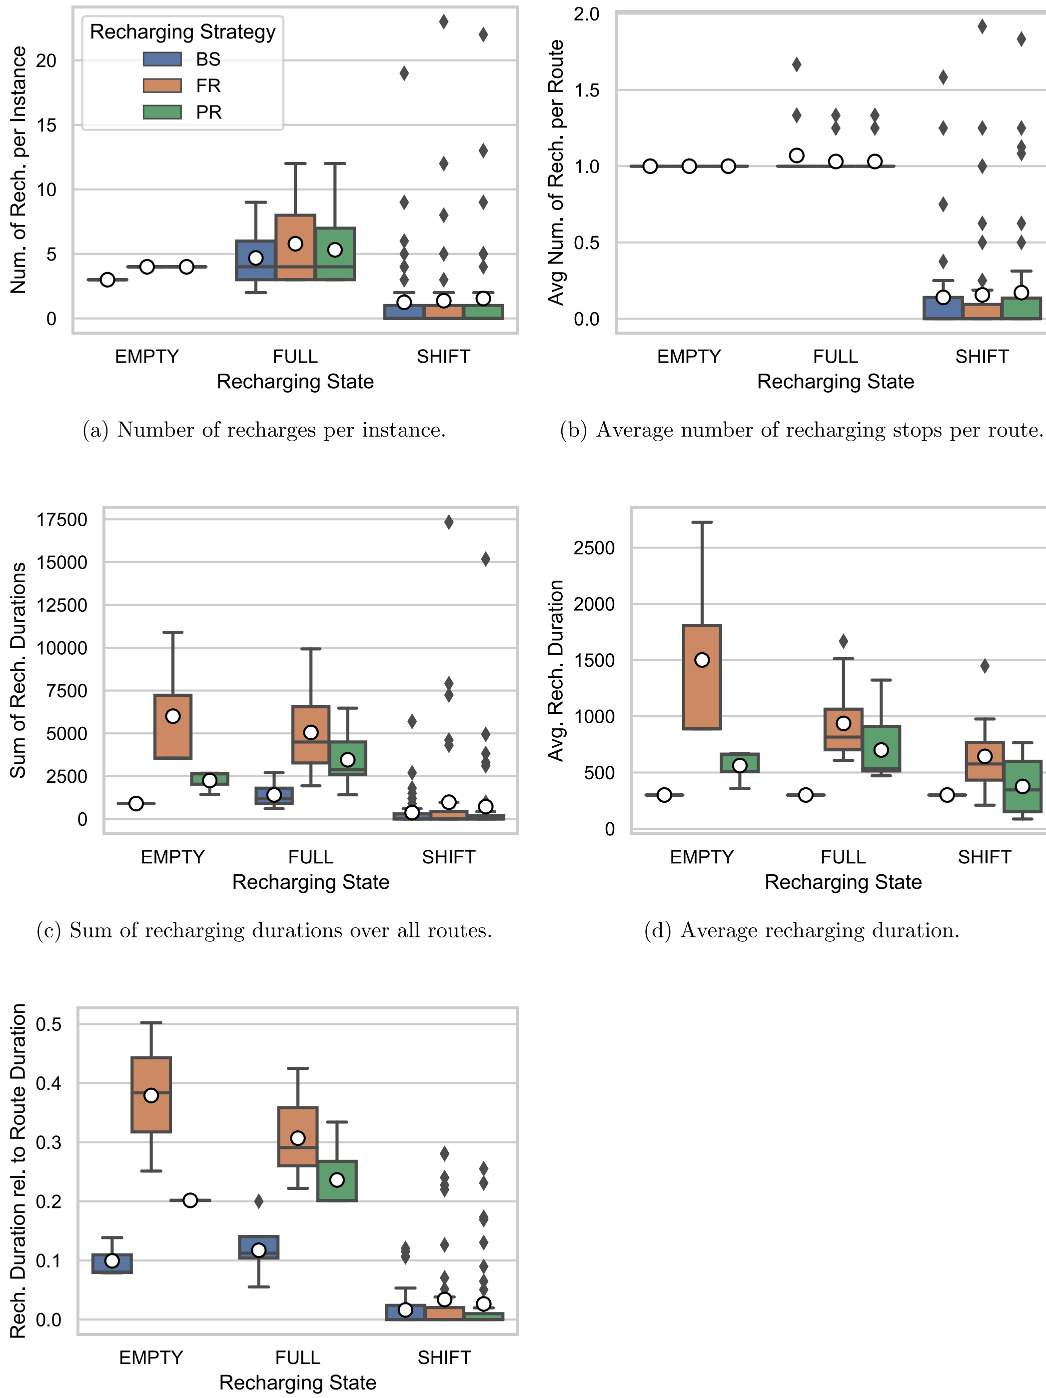

(a) Number of recharges per instance.

(b) Average number of recharging stops per route.

(c) Sum of recharging durations over all routes.

(d) Average recharging duration.

(e) Total recharging duration relative to route duration.

Fig. 8. Recharging behavior over 100 instances for which an optimal solution was found for the DI objective by recharging strategies.

is beneficial. This has to be considered when designing new heuristics for the problem.

- In general, we can state that the battery recharging strategy mainly impacts the feasibility and, hence, the fleet size. If a sufficiently large vehicle fleet is in place, the CT objective is considerably more sensitive to the recharging strategy than the DI objective. This means that if time is critical, e.g., in the manufacturing context, PR or BS are advantageous. Our results also indicate that it depends on the setup whether BS or PR performs better. Hence, the use of case-specific assessments considering also the investment cost is necessary.

7. Conclusions and outlook

In this paper, we investigated the problem of AMR route planning in an intralogistics context. To this end, we introduced several variants of an electric vehicle routing problem with due dates considering different

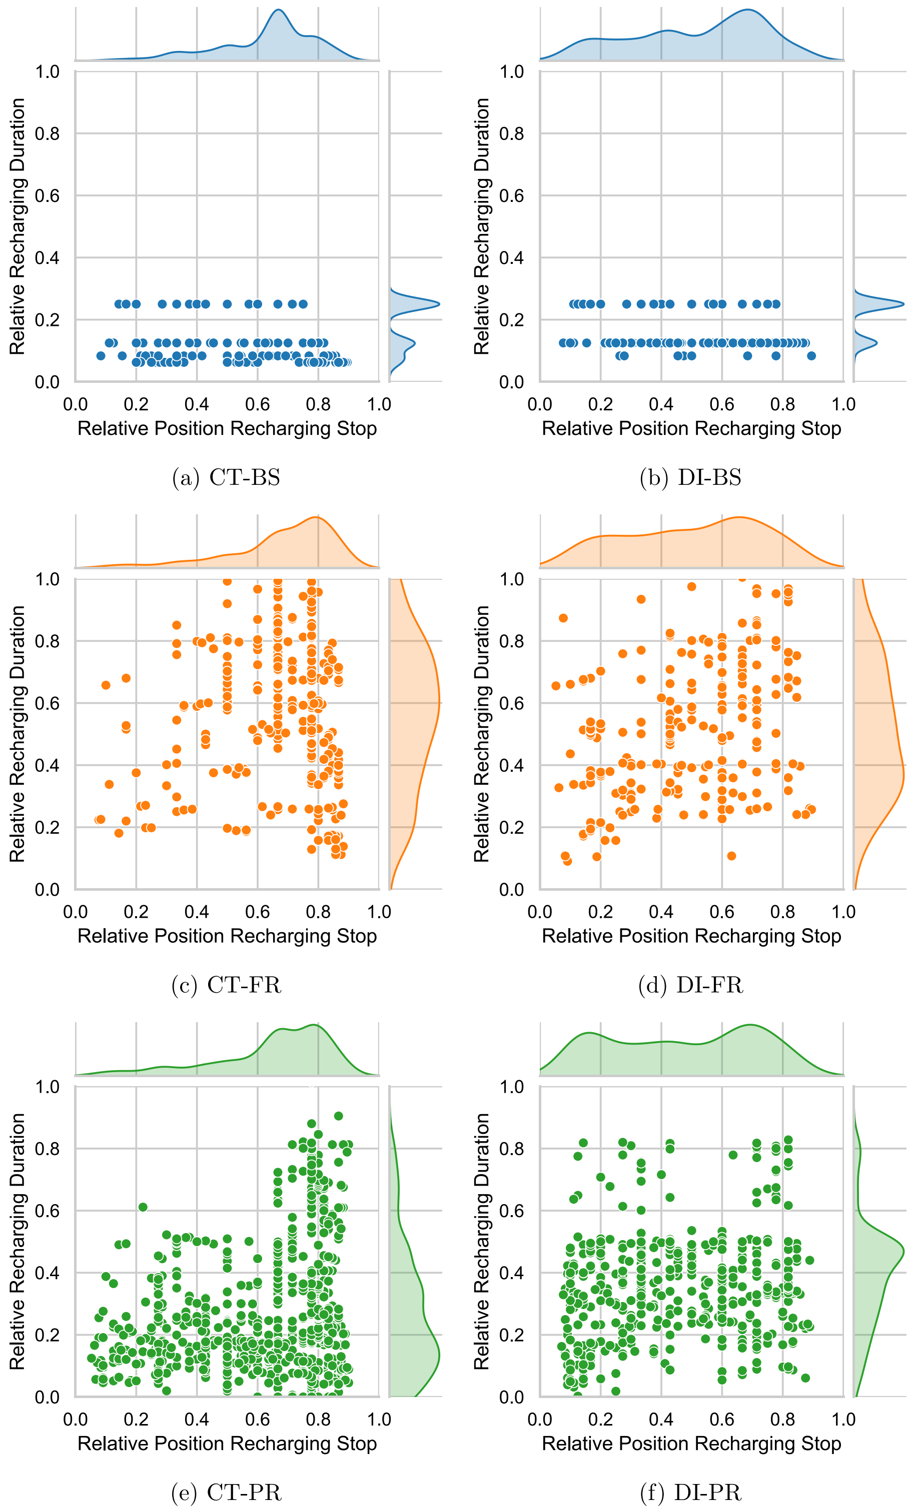

Fig. 9. Relative recharging durations and relative positions of recharging stops over all routes of all optimally solved instances by recharging strategy and objective.

recharging strategies and objective functions. For recharging, we studied battery swapping (BS) and inductive recharging with charging mats (IR). For the latter, we further considered the widely applied recharging strategy with full recharges (FR) and the considerably more complex recharging strategy with partial recharges (PR). As objective functions, the classical minimization of travel distances and the minimization of the total completion time of all jobs were investigated, where the latter is of particular interest in the intralogistics context.

We proposed an effective branch-price-and-cut approach (BPC) for solving all six problem variants. From an algorithmic point of view, the core contribution of this paper is the development of an ad-hoc labeling algorithm for solving the pricing problems of the PR variant,

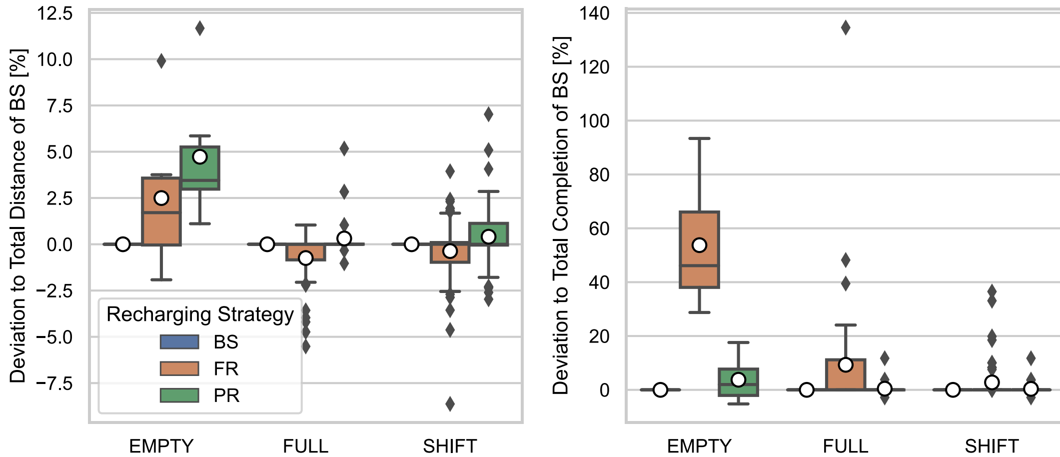

(a) Percentage deviation to total distance. (b) Percentage deviation to total completion time.

Fig. 10. Percentage deviations of BS solutions over 100 instances for which an optimal solution was found for the CT objective by recharging strategies.

minimizing the total completion time. The crucial point here is dealing with the tradeoffs between battery level and time as well as battery level and cost (in terms of total completion time of all jobs).

We reported an extensive computational study on instances adapted from the Solomon standard benchmark set (`Sol`) and a newly introduced benchmark set from an industrial intralogistics use case (`Karis`) to demonstrate the performance of our exact algorithms. In the case of the `Sol` dataset, we analyzed 1698 instances with up to 50 customers. Across all recharging strategies, the BPC approach consistently solved a slightly higher proportion of instances to optimality under the DI objective (76%–79%) compared to the CT objective (66%–69%) within a run time limit of 3600 s. In the case of the `Karis` dataset, we evaluated 1584 instances with 12 to 144 customers. We solved 74% of the instances, proving either optimality or infeasibility for the DI objective and 63% for the CT objective, also with a run time limit of 3600 s.

Apart from the computational aspects, we extensively analyzed the impact of the different recharging strategies and objective functions by comparing solution quality and discussing structures of optimal solutions. Surprisingly, PR shows very good objective values compared to BS, while the simplifying FR strategy seems inappropriate for the addressed context. The results of our study motivate the development of new heuristics for PR, while the typical solution structures should be used in their development.

In future work, broader studies incorporating more instances with different layouts and positioning of recharging stations or transport job situations could be considered to evaluate other logistics scenarios: Aspects such as a capacity larger than one, release dates for transport jobs, a fleet size varying over time, or dynamically arriving transport jobs might be relevant. To find a suitable recharging strategy for a specific application context, the investment cost and maintenance for the infrastructure need to be incorporated.

Apart from the intralogistics use case, modern AMRs are also attractive for other domains. Delivery robots are currently considered as one of several concepts to tackle challenges caused by the continuously increasing number of packages in urban areas (Boysen, Fedtke, & Schwerdfeger, 2021; Mangiaracina et al., 2019). In package delivery, "companies like Amazon want to make deliveries to customers as soon as possible" (Poikonen, Wang, & Golden, 2017), meaning that the total completion time as an objective also plays a crucial role in route planning. However, the characteristics of package delivery differ from those of the context of intralogistics. Hence, in future work, we would like to extend our models to urban logistics and consider related benchmark instances. In this context, it could also be relevant to take into account the possibly limited capacities of recharging stations.

CRediT authorship contribution statement

Anne Meyer: Writing – review & editing, Writing – original draft, Visualization, Software, Methodology, Formal analysis, Data curation, Conceptualization. **Timo Gschwind:** Writing – review & editing, Writing – original draft, Visualization, Software, Methodology, Formal analysis, Data curation, Conceptualization. **Boris Amberg:** Writing – review & editing, Writing – original draft, Visualization, Software, Methodology, Formal analysis, Data curation, Conceptualization. **Dominik Colling:** Validation, Resources, Data curation, Conceptualization.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work the authors used ChatGPT, Grammarly, and DeepL in order to improve the readability and language of their own writing. After using these services, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

Appendix A. Transformation of EVRP with pickup and delivery to EVRP

As stated in Section 3.2, the considered AMR route planning problem corresponds to an EVRP with pickup and delivery (EVRP-PD). Because AMRs have a unit load capacity and can only recharge between two jobs, the EVRP-PD can be modeled as an EVRP. In this section, we explain the transformation.

In the EVRP-PD, a transport job consists of transporting some load from a given pickup location to a given delivery location. If the capacity of the vehicles is assumed to be one and recharging with loaded transport units is not allowed, the AMRs immediately have to drive to the delivery station of a job after they have picked up the load at the corresponding pickup station without the possibility to visit any other location in between. Hence, the transport jobs can be treated as visiting a single customer in an EVRP. To this end, the following input data needs to be transformed:

- Each customer $n \in N$ represents a transport job. Visiting a customer corresponds to performing this transport job, i.e., picking up the load at the pickup location, transporting the load from the pickup to the delivery station, and delivering the load at the delivery location.

Table 12
Detailed characteristics of Karis instances.

Instance group	#inst	Num Cust n	Num Veh $\|K\|$	Rch. Sp. h_{ij}/t_{ij}	Battery cap H^{max}	Init battery H^{init}	Final battery H^{final}
Scenario size							
S	96	{12,24,36,48}	4	{0.25,0.5}	{9.75,19.5,29.25,39}	{0.2,0.8} H^{max}	{0.2,0.8} H^{max}
M	60	{24,48,72,96}	8	{0.25,0.5}			
L	60	{36,72,108,144}	12	{0.25,0.5}			
XL	48	{48,96}	16	{0.25,0.5}			
Number of cycles							
1	96	{12,24,36,48}	{4,8,12,16}	{0.25,0.5}	{9.75,19.5,29.25,39}	{0.2,0.8} H^{max}	{0.2,0.8} H^{max}
2	96	{24,48,72,96}	{4,8,12,16}	{0.25,0.5}			
3	36	{36,72}	{4,8,12,16}	{0.25,0.5}			
4	36	{48,96,144}	{4,8,12,16}	{0.25,0.5}			
Recharging speed							
SLOW	396	{12,24,36,48, 72,96,108,144}	{4,8,12,16}	0.25	{9.75,19.5,29.25,39}	{9.75,19.5, 29.25,39}	{9.75,19.5, 29.25,39}
FAST	396	{12,24,36,48, 72,96,108,144}	{4,8,12,16}	0.5			

- Consider an arc $(i,j) \in A$. If $i \in N$, then i geographically represents the delivery station of the corresponding job. If $j \in N$, then j geographically represents the pickup station of the corresponding job.
- If $i \in N$, the parameter t_{ij} represents the time for performing the job corresponding to customer $i \in N$ as well as the time to reach vertex $j \in J$. More precisely, t_{ij} then comprises the service time at the pickup location, the travel time from the pickup location to the delivery location, and the service time at the delivery location of the transport job corresponding to customer $i \in N$. Furthermore, it comprises the travel time from the corresponding delivery location of i to the geographic location represented by j, i.e., the pickup location corresponding to j if $j \in N$, the location of recharging station $j \in R$, or the location of destination depot $j \in D$.
- The interpretation of parameter h_{ij} with $i \in N$ is analog to t_{ij}. It includes the battery consumption for handling the load at the pickup location, traveling with load from the pickup to the delivery location, and for handling the load at the pickup location of the job corresponding to customer i. Additionally, it includes the battery consumption for traveling without load from the delivery location corresponding to i to the geographic location represented by j.
- Similarly, if $i \in N$, then parameter c_{ij} comprises the travel distance from pickup to delivery location of the job corresponding to customer i as well as the travel distance from the delivery location corresponding to i to the geographic location represented by j.

Appendix B. Detailed characteristics of Karis instances

Table 12 gives an overview of the instance characteristics differentiated by the three instance groups Scenario Size, Number of Cycles, and Recharging Speed.

Appendix C. Detailed computational results for the Sol instances

In this section, we discuss the structure of the resulting routes, applying the different objectives and recharging strategies in more depth. As in Section 5.2.2, we filter the results to the 133 instances for which we solved all six variants – two objectives with three recharging strategies – to optimality.

Fig. 11(a) displays the share of available vehicles that are used in the optimal solutions. CT solutions consistently use always all vehicles in all recharging strategies, while in DI solutions, the number of used vehicles corresponds to the number of vehicles determined in the preprocessing step. This is because the DI objective incentivizes long but fewer routes, avoiding the extra distances from and to the origin and destination depot of the vehicles and increasing the probability of good customer combinations. The CT objective, on the other hand, benefits from scheduling all available vehicles. The positive effect for the completion times results from shorter routes with earlier completion times of the contained customer visits, and from the usage of the initial energy of the whole vehicle fleet.

Fig. 11(b) shows that, on average, the number of recharges in CT solutions is higher than in DI solutions. Apparently, the DI objective tries to avoid frequent recharges to reduce detours to the recharging stations. In contrast, the CT objective accepts more and preferably short recharging stops to avoid the postponement of customer visits later on the routes. This difference is most evident for PR where the durations of the recharging stops can be kept as short as possible.

With the DI objective, the selection of a recharging station strongly depends on the shortest detour. Furthermore, we can observe a tendency to recharge in the middle of the route which can be explained as follows. Visiting a recharging station early on the route is only a good choice if the delays of the arrivals at later stops have no negative effects on the due dates. Visiting a recharging station late on the route is only possible if the battery level allows it. The charts in Fig. 12 illustrate that this effect consistently occurs for all recharging strategies. On the x-axis, they show the positions of the recharge stops relative to all stops of a route for all routes of all solutions. For the DI objective, the distribution is quite symmetric, with a mean of 0.5 and a rather high variance. Fig. 11(c) shows that less than 20% of the routes contain a visit to a recharging station just before returning to the destination depot. On the contrary, in CT solutions, the share of routes for which the last vertex before the destination depot is a recharging station is high for all recharging strategies. For PR, this happens in almost all routes. Accordingly, in Fig. 12, the distributions of the relative positions of the recharging stops are strongly skewed to the left for all recharging strategies.

To analyze the impact of the objective function and the recharging strategies on the length of the recharging stops, Fig. 12 depicts on the y-axis the recharging duration per recharging stop relative to the planning horizon of the corresponding instance. Recall that for BS, the battery swap duration was set to 10% of the battery capacity (H_v^{max}) and, hence, the duration of each recharging stop of a given instance is constant. The differences in the chart result only from different battery capacities of the instances and should not be interpreted otherwise.

For FR, the recharging duration depends directly on the current battery charge level, and hence, a recharging stop relatively early in

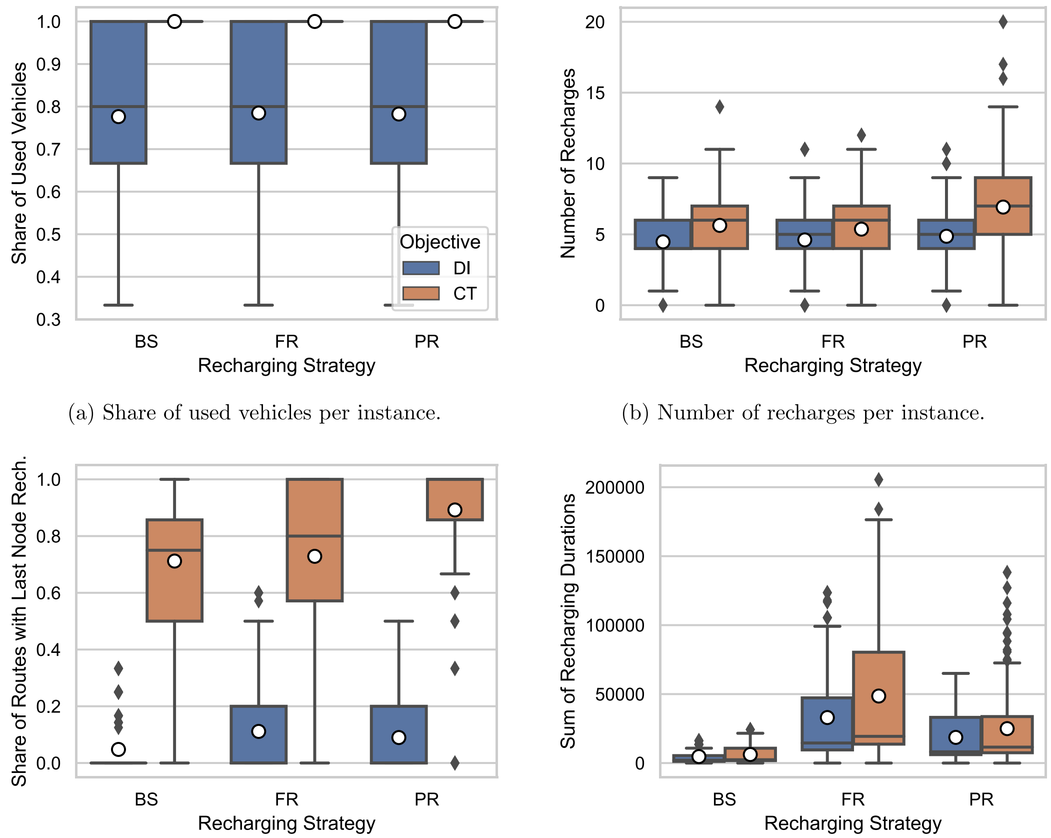

(a) Share of used vehicles per instance.

(b) Number of recharges per instance.

(c) Share of routes for which the last vertex before the destination depot is a recharging station.

(d) Sum of recharging durations over all routes per instance.

Fig. 11. Recharging behavior over 133 instances for which an optimal solution was found for all six variants.

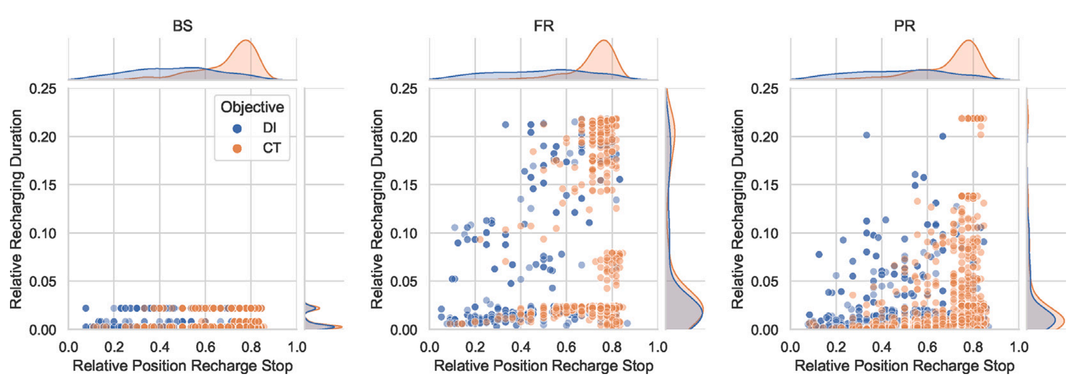

Fig. 12. Relative recharging duration and relative position of recharge stop over all routes of all instances by recharging strategy (different plots) and objective (color). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

the route results in relatively short recharging durations. Even more, Fig. 12 shows that also most of the late recharging stops still feature rather short recharging durations, indicating consistently high battery charge levels. For the CT objective, however, a considerable share of the late recharging stops take around 20% of the planning horizon. This explains the difference in the sums of the recharging durations between DI and CT depicted in Fig. 11(d).

For PR, Fig. 12 reveals that most recharging stops are rather short, irrespective of the objective function. This shows that the additional degree of freedom of this strategy is heavily used, especially with the CT objective. The strongly left-skewed distribution of the relative recharging duration in Fig. 12 and the number of recharges per instance in Fig. 11(b) mean that solutions for the variant with CT and PR comprise comparatively many but short recharging stops. This is especially noteworthy because the use of the whole vehicle fleet leads to higher initial battery charge levels, and yet more recharge stops are scheduled.

References

Baldacci, R., Mingozzi, A., & Roberti, R. (2011). New route relaxation and pricing strategies for the vehicle routing problem. *Operations Research*, *59*(5), 1269–1283.

Boysen, N., Fedtke, S., & Schwerdtfeger, S. (2021). Last-mile delivery concepts: a survey from an operational research perspective. *OR Spectrum*, *43*(1), 1–58.

De Ryck, M., Versteyhe, M., & Debrouwere, F. (2020). Automated guided vehicle systems, state-of-the-art control algorithms and techniques. *Journal of Manufacturing Systems*, *54*, 152–173.

De Ryck, M., Versteyhe, M., & Shariatmadar, K. (2020). Resource management in decentralized industrial automated guided vehicle systems. *Journal of Manufacturing Systems*, *54*, 204–214.

Desaulniers, G., Desrosiers, J., & Solomon, M. M. (2002). Accelerating strategies in column generation methods for vehicle routing and crew scheduling problems. In *Essays and surveys in metaheuristics* (pp. 309–324). Boston, MA: Springer US.

Desaulniers, G., Errico, F., Irnich, S., & Schneider, M. (2016). Exact algorithms for electric vehicle-routing problems with time windows. *Operations Research*, *64*(6), 1388–1405.

Desaulniers, G., Gschwind, T., & Irnich, S. (2020). Variable fixing for two-arc sequences in branch-price-and-cut algorithms on path-based models. *Transportation Science*, *54*(5), 1170–1188.

Duman, E. N., Taş, D., & Çatay, B. (2022). Branch-and-price-and-cut methods for the electric vehicle routing problem with time windows. *International Journal of Production Research*, *60*(17), 5332–5353.

Ebben, M. (2001). *Logistic control in automated transportation networks* (Ph.D. thesis), University of Twente.

Erdelić, T., & Carić, T. (2019). A survey on the electric vehicle routing problem: Variants and solution approaches. *Journal of Advanced Transportation*, *2019*, 1–48.

Fragapane, G., de Koster, R., Sgarbossa, F., & Strandhagen, J. O. (2021). Planning and control of autonomous mobile robots for intralogistics: Literature review and research agenda. *European Journal of Operational Research*, *294*(2), 405–426.

Furmans, K., Seibold, Z., & Trenkle, A. (2019). Future technologies in intralogistics and material handling. In *Operations, logistics and supply chain management* (pp. 545–574). Springer.

GEBHARDT Intralogistics Group (2022). GEBHARDT KARIS® Autonomes fahrerloses Transportsystem. https://www.gebhardt-foerdertechnik.de/de/produkte/fahrerlose-transportsysteme/gebhardt-karis/. (Accessed 04 November 2022).

Gschwind, T., Irnich, S., Tilk, C., & Emde, S. (2019). Branch-cut-and-price for scheduling deliveries with time windows in a direct shipping network. *Journal of Scheduling*, *23*(3), 363–377.

Hiermann, G., Puchinger, J., Ropke, S., & Hartl, R. F. (2016). The electric fleet size and mix vehicle routing problem with time windows and recharging stations. *European Journal of Operational Research*, *252*(3), 995–1018.

Irnich, S. (2007). Resource extension functions: properties, inversion, and generalization to segments. *OR Spectrum*, *30*(1), 113–148.

Irnich, S., & Desaulniers, G. (2005). Shortest path problems with resource constraints. In G. Desaulniers, J. Desrosiers, & M. Solomon (Eds.), *Column generation* (pp. 33–65). Springer.

Jepsen, M., Petersen, B., Spoorendonk, S., & Pisinger, D. (2008). Subset-row inequalities applied to the vehicle-routing problem with time windows. *Operations Research*, *56*(2), 497–511.

Jun, S., Lee, S., & Yih, Y. (2021). Pickup and delivery problem with recharging for material handling systems utilising autonomous mobile robots. *European Journal of Operational Research*, *289*(3), 1153–1168.

Kabir, Q. S., & Suzuki, Y. (2018). Increasing manufacturing flexibility through battery management of automated guided vehicles. *Computers & Industrial Engineering*, *117*, 225–236.

Kabir, Q. S., & Suzuki, Y. (2019). Comparative analysis of different routing heuristics for the battery management of automated guided vehicles. *International Journal of Production Research*, *57*(2), 624–641.

KARIS PRO Consortium (2016). Final Project Report: KARIS PRO – Autonomer Materialtransport für flexible Intralogistik. http://karispro.de/Abschlussbericht KARIS PRO.pdf.

Keskin, M., & Çatay, B. (2016). Partial recharge strategies for the electric vehicle routing problem with time windows. *Transportation Research Part C (Emerging Technologies)*, *65*, 111–127.

Keskin, M., Laporte, G., & Çatay, B. (2019). Electric vehicle routing problem with time-dependent waiting times at recharging stations. *Computers & Operations Research*, *107*, 77–94.

Kohl, N., Desrosiers, J., Madsen, O., Solomon, M., & Soumis, F. (1999). 2-Path cuts for the vehicle routing problem with time windows. *Transportaton Science*, *33*(1), 101–116.

Kucukoglu, I., Dewil, R., & Cattrysse, D. (2021). The electric vehicle routing problem and its variations: A literature review. *Computers & Industrial Engineering*, *161*, Article 107650.

Lam, E., Desaulniers, G., & Stuckey, P. J. (2022). Branch-and-cut-and-price for the electric vehicle routing problem with time windows, piecewise-linear recharging and capacitated recharging stations. *Computers & Operations Research*, *145*, Article 105870.

Le-Anh, T., & De Koster, M. (2006). A review of design and control of automated guided vehicle systems. *European Journal of Operational Research*, *171*(1), 1–23.

Lera-Romero, G., Bront, J. J. M., & Soulignac, F. J. (2024). A branch-cut-and-price algorithm for the time-dependent electric vehicle routing problem with time windows. *European Journal of Operational Research*, *312*(3), 978–995.

Mangiaracina, R., Perego, A., Seghezzi, A., & Tumino, A. (2019). Innovative solutions to increase last-mile delivery efficiency in B2c e-commerce: a literature review. *International Journal of Physical Distribution & Logistics Management*, *49*(9), 901–920.

McHaney, R. (1995). Modelling battery constraints in discrete event automated guided vehicle simulations. *International Journal of Production Research*, *33*(11), 3023–3040.

Meyer, A., Gschwind, T., Amberg, B., & Colling, D. (2024). BEAMI - Benchmark instances for routing electric autonomous mobile robots in intralogistics (version V1). Zenodo dataset: https://doi.org/10.5281/zenodo.14284157.

Montoya, A., Guéret, C., Mendoza, J. E., & Villegas, J. G. (2017). The electric vehicle routing problem with nonlinear charging function. *Transportation Research, Part B (Methodological)*, *103*, 87–110.

Naddef, D., & Rinaldi, G. (2002). Branch-and-Cut algorithms for the capacitated VRP. In P. Toth, & D. Vigo (Eds.), *The vehicle routing problem* (pp. 53–84). Philadelphia: Siam.

Pecin, D., Contardo, C., Desaulniers, G., & Uchoa, E. (2017). New enhancements for the exact solution of the vehicle routing problem with time windows. *INFORMS Journal on Computing*, *29*(3), 489–502.

Pecin, D., Pessoa, A., Poggi, M., & Uchoa, E. (2017). Improved branch-cut-and-price for capacitated vehicle routing. *Mathematical Programming Computation*, *9*(1), 61–100.

Pinedo, M. (2009). *Scheduling theory, algorithms, and systems* (3rd ed.). Springer.

Poikonen, S., Wang, X., & Golden, B. (2017). The vehicle routing problem with drones: Extended models and connections. *Networks*, *70*(1), 34–43.

Schneider, M., Stenger, A., & Goeke, D. (2014). The electric vehicle-routing problem with time windows and recharging stations. *Transportation Science*, *48*(4), 500–520.

Schulz, A. (2024). Using infeasible path cuts to solve electric vehicle routing problems with realistic charging functions exactly within a branch-and-cut framework. *EURO Journal on Transportation and Logistics*, *13*, 100–131.

Selmair, M., Hauers, S., & Gustafsson-Ende, L. (2019). Scheduling charging operations of autonomous AGVs in automotive in-house logistics. In M. Putz, & A. Schlegel (Eds.), *simulation in produktion und logistik 2019* (pp. 315–324).

Solomon, M. M. (1987). Algorithms for the vehicle routing and scheduling problems with time window constraints. *Operations Research*, *35*(2), 254–265.

Vis, I. F. (2006). Survey of research in the design and control of automated guided vehicle systems. *European Journal of Operational Research*, *170*(3), 677–709.

Wang, W., Adulyasak, Y., Cordeau, J., & He, G. (2024). The heterogeneous-fleet electric vehicle routing problem with nonlinear charging functions. *Les Cahiers du GERAD ISSN*, *711*, 2440.

Zou, B., Xu, X., De Koster, R., et al. (2018). Evaluating battery charging and swapping strategies in a robotic mobile fulfillment system. *European Journal of Operational Research*, *267*(2), 733–753.