

Joint Routing of Conventional and Range-Extended Electric Vehicles in a Large Metropolitan Network

Anirudh Subramanyam^{*1}, Taner Cokyasar^{1,2}, Jeffrey Larson¹, and Monique Stinson¹

¹Argonne National Laboratory, 9700 S. Cass Ave., Lemont, IL 60439, USA

²Tarsus University, Takbas mah., Kartaltepe sk., Tarsus, Mersin 33400, Turkey

January 4, 2022

Abstract

Range-extended electric vehicles combine the higher efficiency and environmental benefits of battery-powered electric motors with the longer mileage and autonomy of conventional internal combustion engines. This combination is particularly advantageous for time-constrained delivery routing in dense urban areas, where battery recharging along routes can be too time-consuming to economically justify the use of all-electric vehicles. However, switching from electric to conventional fossil fuel modes also results in higher costs and emissions and lower efficiency. This paper analyzes this heterogeneous vehicle routing problem and describes two solution methods: an exact branch-price-and-cut algorithm and an iterated tabu search metaheuristic. From a methodological perspective, we find that the exact algorithm consistently obtains tight lower bounds that also serve to certify the metaheuristic solutions as near-optimal. From a policy standpoint, we examine a large-scale real-world case study concerning parcel deliveries in the Chicago metropolitan area and quantify various operational metrics including energy costs and vehicle miles traveled. We find that by deploying roughly 20% of range-extended vehicles with a modest all-electric range of 33 miles, parcel distributors can save energy costs by up to 17% while incurring less than 0.5% increase in vehicle miles traveled. Increasing the range to 60 miles further reduces costs by only 4%, which can alternatively be achieved by decreasing the average service time by 1 minute or increasing driver working time by 1 hour. Our study reveals several key areas of improvement on which vehicle manufacturers, distributors, and policy makers can focus their attention.

Keywords: Large-scale vehicle routing, electric vehicles, delivery planning, optimization.

1 Introduction

The vehicle routing problem (VRP) and its many variants have been widely studied in the field of transportation science and logistics. Yet, the inherent combinatorial complexity of such problems, along with the emergence of new vehicles with distinct characteristics—such as electric vehicles (EVs), plug-in hybrid EVs (PHEVs), autonomous vehicles, and unmanned aerial vehicles—make VRPs challenging to solve. A common characteristic underlying virtually all VRP variants is the determination of optimal routes for a fleet of vehicles to serve transportation requests distributed across geographic locations. Although the emergence of new vehicle technologies does not change this requirement, such technologies can introduce additional complicating constraints that require new modeling approaches.

This paper is motivated by the growing federal and private impetus [The White House, 2021, Bogage, 2021] for the adoption of range-extended EVs (REEVs) and PHEVs as delivery trucks for urban transportation, including e-commerce and postal operations. REEVs are hybrid vehicles that extend the limited driving range of battery-electric vehicles (BEVs) using a fuel-based power unit (i.e., an internal combustion engine). This extended range can be achieved either via an electric generator that charges the battery once it reaches a minimum charge state (“charge-sustaining” mode) after the vehicle has exhausted its all-electric range

^{*}Corresponding author: asubramanyam@anl.gov

(“charge-depleting” mode) or by constantly supplementing the battery energy (“blended” mode). Irrespective of the mode, the crucial point is that because the batteries do not require external recharging, REEVs combine the efficiency and environmental benefits of BEVs with the range and autonomy of conventional vehicles (CVs). That is, REEVs can alleviate the “range anxiety” that can accompany pure BEVs.

Our motivation stems from the potential use of REEVs for distribution logistics. The Clean Energy Research Center - Truck Research Utilizing Collaborative Knowledge (CERC TRUCK) is a U.S. Department of Energy-sponsored consortium between industry and academia that is developing a compressed natural gas (CNG) range extender medium-duty vehicle for the pickup-and-delivery market, with the goal of increasing freight efficiency by 50% over a 2016 baseline [U.S.-China Clean Energy Research Center, 2021]. Other CNG REEVs will be on the roads shortly: a new Class-6 REEV delivery truck has been developed by Efficient Drivetrains, Inc., that runs on a CNG-series hybrid-electric powertrain [Ferchau, 2017, Guanetti et al., 2016, Jayanthan, 2017], and a Class-8 CNG-powered hybrid-electric commercial truck has been developed by Hyliion Inc. with planned volume shipments in 2022 [Green Car Congress, 2020]. Considering that CNG is environmentally cleaner and cheaper than both gasoline and diesel [Yadav et al., 2019, Goreham, 2014] but has an equally robust refueling infrastructure, replacing conventional parcel delivery trucks with such CNG-based REEVs can simultaneously decrease environmental impacts and operational costs. In fact, these reductions may be achieved by using diesel-based REEVs as well; for example, it has been shown that hybrid conversion of light-duty diesel delivery trucks can already improve fuel economy by more than 24% [Byun and Choi, 2021]. While much of the discussion in this manuscript centers on CNG-based REEVs, the developed methods are naturally applicable to REEVs independent of their conventional fuel source.

In this study we investigate the impact of REEVs on large-scale delivery operations in terms of multiple performance metrics, including energy cost, fleet size, vehicle miles traveled (VMT), and vehicle hours traveled (VHT). We consider various realistic scenarios where a mix of CVs and REEVs are available for delivery operations. We define and solve a new VRP variant in these scenarios that simultaneously determines the optimal fleet composition and truck routes under various operational constraints, including limited EV ranges, truck capacities, and driver hours-of-work limits. This range-extended electric vehicle routing problem (REEVRP) is an instance of the heterogeneous VRP (HVRP), where the goal is to determine the optimal fleet composition and vehicle routes for a heterogeneous fleet of vehicles with different capacities and costs. Although the REEVRP also features a heterogeneous fleet of CVs and REEVs, what distinguishes the REEVRP from a classical HVRP is the heterogeneity within the REEV vehicle type. That is, REEVs feature two distinct propulsion modes: a conventional fuel mode and a (limited-range) battery-powered EV mode.¹ The cost of the EV mode (i.e., the electricity cost to charge batteries) is considerably lower than that of the conventional fuel mode (i.e., the cost of CNG or gasoline). Consequently, the total cost of an REEV route is no longer a linear function of the corresponding travel time or distance, and modeling REEVs as a single vehicle type within a classical HVRP model can lead to a poor approximation of the actual costs. In addition to this nonlinear cost structure, the REEVRP we consider captures other realistic constraints, including multiple vehicle types, vehicle capacities, driving time limits, and asymmetric travel times and vehicle speeds.

The contributions of this paper are twofold. First, from a methodological standpoint, we formally define the REEVRP as a new VRP variant and propose two algorithms for its solution: (i) an exact branch-price-and-cut (BPC) algorithm based on a novel set partitioning integer programming (IP) formulation and (ii) an iterated tabu search (ITS) metaheuristic that can provide solutions for large-scale problem instances with more than 3,000 delivery locations. The BPC algorithm consistently certifies the ITS solutions as near-optimal for small- to medium-scale instances with up to 100 locations. Second, from a practical standpoint, we apply these methods to analyze the impacts of REEVs on the Chicago metropolitan area under various realistic operational scenarios where we vary the proportions of CVs and EVs, EV ranges, hours-of-work limits, service times, and vehicle capacities. We provide important managerial insights by quantifying systemwide performance metrics under these scenarios to identify the key parameters that can guide policy and decision-making. We find the following:

- Distributors can expect energy cost savings of up to 17% while incurring less than a 0.5% increase in VMT and VHT by deploying a small number of REEVs per depot (roughly 20% of the fleet) each with

¹We use the term “EV mode” to refer to either the charge-depleting mode or the mileage-equivalent blended mode, depending on the hybrid vehicle technology.

a modest EV range of 33 miles.

- The marginal cost savings diminish rapidly with increasing EV range; for example, doubling the range to 60 miles further reduces costs by only 4%. The reason is partly because delivery routes are time-constrained. Indeed, decreasing the average service time by 1 minute and increasing driver working time by 1 hour can decrease both the energy costs and VMT by roughly 8% and 4%, respectively. For the same reason, increasing vehicle capacity is unlikely to provide cost or VMT reductions (for the scenarios we consider).

These findings along with key operational statistics reveal potential areas of improvement where vehicle manufacturers, logistics service providers, and policy makers can focus their attention.

The outline of the manuscript is as follows. [Section 2](#) provides a comprehensive literature review of similar studies from a methodological perspective. [Section 3](#) describes the REEVRP setting, provides underlying assumptions, and briefly mentions data sources utilized. [Section 4](#) formulates a mathematical IP model to solve REEVRP instances and describes an exact BPC method. [Section 5](#) illustrates the ITS metaheuristic for solving large-scale REEVRP instances. [Section 6](#) provides case studies of the computational performance of the solution methodologies; these studies also provide managerial insights for various deployment scenarios. [Section 7](#) concludes with a discussion of the strengths and weaknesses of our study and provides potential future research directions.

2 Literature Review

The VRP literature is rich; see [Toth and Vigo \[2014\]](#) for a broad overview. The REEVRP belongs to the subclass of VRPs featuring multiple vehicle types, collectively known as heterogeneous VRPs, of which comprehensive reviews can be found in [Hoff et al. \[2010\]](#), [Irnick et al. \[2014\]](#), [Koç et al. \[2016\]](#). A distinguishing feature of the REEVRP is the use of hybrid or electric vehicle types. VRP variants that incorporate EVs have burgeoned over the past decade; see [Pelletier et al. \[2016\]](#), [Erdelić and Carić \[2019\]](#) for recent surveys. For brevity, we discuss only similarities and differences of our work with the most directly relevant literature.

The unique features of the relevant EV-related works in the literature are the limited battery capacities and driving ranges of pure BEVs. This limitation requires optimizing over decisions corresponding to recharging station visits in addition to standard routing decisions. Seminal papers that first considered this additional decision layer include those of [Conrad and Figliozzi \[2011\]](#), [Erdoğan and Miller-Hooks \[2012\]](#) and [Schneider et al. \[2014\]](#). Since then, various heuristic and exact solution approaches have been proposed to more realistically model the real-world routing of EVs, including more realistic charging models (e.g., partial recharging) or models where the battery charge state is a nonlinear function of the charging time [[Felipe et al., 2014](#), [Montoya et al., 2016](#), [Desaulniers et al., 2016](#), [Keskin and Çatay, 2016](#), [Froger et al., 2019](#)]. Extensions to heterogeneous fleets of BEVs and CVs have also been studied [[Juan et al., 2014](#), [Sassi et al., 2015](#), [Goeke and Schneider, 2015](#), [Lebeau et al., 2015](#), [Hiermann et al., 2016](#)].

The distinguishing aspect of our work from pure BEV routing problems is the use of range extenders to obviate the need to plan for time-consuming visits to recharging stations. Pure BEV routing problems where battery recharging is disallowed en route, such as in dense urban environments, have been studied by [Pelletier et al. \[2018\]](#) and [Florio et al. \[2021\]](#). These papers consider homogeneous fleets of BEVs and include detailed models for battery charge scheduling (at the depot) and also ensure that battery capacity is not exceeded during delivery operations. However, the absence of alternative propulsion modes in the same vehicle leads to a simple linear cost function for each vehicle, be it the total energy cost or a distance-proportional travel cost.

The literature on REEV and PHEV routing is relatively sparse. One of the earliest works is that of [Abdallah \[2013\]](#), where the problem is to minimize the cost of using conventional fuels in a homogeneous fleet of PHEVs that are used to serve customers in predefined time windows. A similar problem is considered by [Mancini \[2017\]](#), where the author presents a heuristic approach to route a homogeneous fleet of PHEVs that can switch between EV and conventional fuel modes at any point in the route. These models were generalized by [Hiermann et al. \[2019\]](#), who consider a heterogeneous fleet composed of pure BEVs, PHEVs, and CVs, with the objective of finding the optimal fleet composition that minimizes the sum of fixed cost (per used vehicle) and energy costs (for electricity and conventional fuels). All of these studies allow for

battery recharging decisions en route. Whereas Abdallah [2013] allows PHEVs the option to be recharged at potentially every delivery location, Hiermann et al. [2019] and Mancini [2017] allow PHEVs to be recharged only at specific recharging stations. Vincent et al. [2017] and Li et al. [2020] consider a similar variant with homogeneous PHEVs without customer time windows but allow PHEVs to also refuel gasoline (or other conventional fuel) at specific stations. All of these papers propose metaheuristic solution approaches that are evaluated on synthetic test cases. In particular, Abdallah [2013] proposes a tabu search approach; Hiermann et al. [2019] design a sophisticated layered metaheuristic consisting of a genetic algorithm and large neighborhood search hybridized with integer programming; Mancini [2017] design a large neighborhood search metaheuristic; and Vincent et al. [2017] and Li et al. [2020] propose simulated annealing and variable neighborhood descent methods, respectively. These studies consider instances of up to 100 locations and report solution times on the order of a few hours [Abdallah, 2013], 10 minutes [Mancini, 2017, Hiermann et al., 2019], and 1 minute [Vincent et al., 2017, Li et al., 2020].

In contrast to the aforementioned papers, our work does not consider en route battery recharging decisions. In other words, we assume that all vehicles charge only at the depot and that they begin each route with a full charge. This assumption allows us to design exact methods that can quickly solve problems with up to 100 locations and scalable metaheuristic approaches that can approximately solve real-world test cases consisting of more than 3,000 locations in less than an hour. Along these lines, problems involving hybrid electric vehicles that do not involve recharging decisions have also been studied by Doppstadt et al. [2016]. However, they consider routing only a single vehicle that must decide which of four different propulsion modes to use on every single arc of its route, with each mode associated with a different cost on each arc. In a similar vein, models that account for the optimal power management of EV and conventional fuel modes—as a function of the travel speed, power requirements, or other arc-level parameters—with varying levels of complexity can be found in the work of Bahrami et al. [2020], De Nunzio et al. [2021], and Casparia et al. [2021]. Although more accurate, these models can be complex (involving mixed-integer nonlinear optimization or nested dynamic programs), often leading to poor scalability on real-world problems. This level of detail cannot address the scale of networks we consider in our study; instead, we make simplifying assumptions about the cost of using alternative propulsion modes that allow the solution of real-scale instances without losing key managerial and policy insights.

3 Problem Description

A growing number of logistics service providers (LSPs) are considering the adoption of REEVs into their distribution fleets to reduce their energy costs and environmental footprint. LSPs are particularly interested in understanding the extent of cost savings (if any) under various deployment scenarios. The magnitude of such savings could be quantified with a reasonably accurate VRP model that captures the essential real-world features.

To that end, we focus our attention on the Chicago metropolitan area, where we assume that four large LSPs—Amazon, FedEx, UPS, and USPS—are tasked with delivering packages on a typical day. We acquire demand and road network data from POLARIS, the Planning and Operations Language for Agent-based Regional Integrated Simulation, developed at Argonne National Laboratory [Auld et al., 2016]. We also identify 53 depot locations of the four LSPs from publicly available sources. We assume the LSPs uniformly share the total number of delivery tasks between them (totaling roughly 600,000 customer locations) and that they each solve an assignment problem to determine their depot-to-customer allocations. Figure 1 illustrates the resulting regional-level problem layout. Further details of the case study and experimental design are deferred to Section 6.

With the regional-level operation decomposed into depot-level operations, we then define and solve the proposed REEVRP for each depot. Although one can also define and solve a multidepot variant of the REEVRP for each LSP (where the depot-to-customer assignment is part of the optimization), the scale of real-world problems requires a decomposition approach, and one can view the REEVRP as a subproblem as part of this overall decomposition strategy. Moreover, even if we focus our attention on a single depot, the considered area in Figure 1 already leads to fairly large REEVRP instances.

The goal of each depot-level REEVRP is to solve a combined fleet composition and routing problem. Each depot must determine the optimal delivery routes for its trucks while simultaneously choosing a truck




Figure 1: Illustrative problem layout in the Chicago metropolitan area. Customer locations are color-coded to match the colors of their assigned depots. Although certain colors appear to dominate the main figure, this is due to overlapping points. The inset shows that customers of all LSPs appear throughout the metropolitan area.

type (either CV or REEV) for each route. The number of delivery trucks of each type is limited. All vehicles are assumed to leave their depot at the beginning of the day, visit customer locations to make deliveries, and then return to their original depot. Individual arc lengths and their speeds are provided by POLARIS; these values define the average travel times in the network. Note that we assume the same average speed for all delivery trucks (on a given arc), although our methodology does not require this assumption. The REEV trucks are assumed to leave the depot with fully charged batteries and do not stop mid-route to recharge. We assume that REEVs can optimally switch between CNG and EV modes at any point along their route and this internal power management is decided by some lower-level logic (e.g., see [Bahrami et al. \[2020\]](#), [Chau et al. \[2017\]](#)), which is outside the scope of our study. Crucially, this implies that their EV range is always used to full capacity (assuming that their routes are longer than the EV range). Once the EV range has been traveled (i.e., the battery is depleted), the CNG-powered internal combustion engine (or some other conventional fuel) kicks in and allows the truck to complete the remainder of its service.

The overall objective of each depot-level REEVRP is to minimize the total energy costs. The energy cost of CVs is readily estimated based on their miles per gallon energy efficiency and the average gallon cost of fuel (diesel in our case), REEVs incur two types of energy costs depending on the mode (EV or CNG) by which they are powered. The energy cost in EV mode is estimated by using the average electricity price and electricity consumption in that mode, whereas the energy cost in CNG mode is estimated based on the fuel economy in that mode and the average cost of CNG, similar to CVs. We do not consider fixed costs per vehicle usage (e.g., to account for ownership, maintenance, or infrastructure-related costs) because of the difficulty in accurately estimating these parameters. We note, however, that the proposed optimization methods do not preclude us from considering such fixed costs in addition to the energy-based variable costs (e.g., by simply adding them to the depot-incident arc travel costs).

All delivery trucks are constrained by both time (i.e., work hours) and physical capacity. This means that the total time corresponding to any route (including road travel and service times) cannot exceed a prespecified number of work hours and that the total number of packages delivered along any route is capped by an upper limit. We note that the time spent for parking trucks, unloading packages, and physically making deliveries are all included in the service time. For simplicity, we assume that all trucks have the same physical carrying capacity and work hours; as before, our models can be readily extended to lift these assumptions. In the following sections we describe the problem formally and provide solution methods to a general form

of the problem. We emphasize that the proposed methodology applies to any mixed REEV/CV fleet and is not tied to our motivating example of CNG-powered REEVs.

4 Mathematical Model and Exact Branch-Price-and-Cut Algorithm

We first provide a precise definition of the REEVRP in Section 4.1, then present our IP formulation of the problem in Section 4.2, and describe its solution via a BPC algorithm in Section 4.3. For ease of reference for the reader, Table 1 provides a centralized summary of the important notation that is formally defined throughout the paper.

Table 1: Notation used in the mathematical model.

Set	Description
\mathcal{G}	Directed graph representation of the road network
\mathcal{V}	Set of depot and delivery nodes in \mathcal{G}
\mathcal{A}	Set of directed arcs in \mathcal{G}
\mathcal{V}_C	Set of delivery nodes in \mathcal{G}
\mathcal{K}	Set of vehicle types
\mathcal{K}^+	Set of vehicle types used in set partitioning model
\mathcal{P}	Set of paths in \mathcal{G} that start at the origin depot 0 and end at the destination depot $n+1$
\mathcal{R}	Set of all feasible elementary routes in \mathcal{G}
\mathcal{R}^{ng}	Set of all feasible ng -routes in \mathcal{G}
$NG_i, i \in \mathcal{V}_C$	ng -set of node i
$\mathcal{C}_p, p \in \mathcal{P}$	Set of arcs in the transitive closure of (the subgraph induced by) path p
Parameter	Description
$q_i, i \in \mathcal{V}$	Demand of node i
$s_i, i \in \mathcal{V}$	Service time of node i
$t_{ij}, (i, j) \in \mathcal{A}$	Travel time along arc (i, j)
$d_{ij}, (i, j) \in \mathcal{A}$	Travel distance along arc (i, j)
$m_k, k \in \mathcal{K}$	Number of available vehicles of type k
Q	Demand carrying capacity of each vehicle
T	Maximum working time (hours-of-work) before which vehicles must return to the depot
D_E	EV mode mileage of REEVs (EV range)
c_e	Mileage cost of REEVs in EV mode
c_g	Mileage cost of REEVs in non-EV mode
c_0	Mileage cost of CVs
$\alpha_{ip}, i \in \mathcal{V}, p \in \mathcal{P}$	Total number of times node i is visited on path p
$\beta_{ijp}, (i, j) \in \mathcal{A}, p \in \mathcal{P}$	Total number of times arc (i, j) appears on path p
$\eta_p, p \in \mathcal{P}$	Total demand of path p
$\delta_p, p \in \mathcal{P}$	Total distance of path p
$\tau_p, p \in \mathcal{P}$	Total processing time (sum of travel and service times) of path p
$c_{pk}, p \in \mathcal{P}, k \in \mathcal{K}$	Cost incurred when path p is performed by a vehicle of type k
Decision Variable	Description
$\lambda_{rk}, r \in \mathcal{R}^{ng}, k \in \mathcal{K}^+$	Binary variable indicating if route r is performed by a vehicle of type k

4.1 Problem Definition

The REEVRP is modeled on a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{A})$ with nodes $\mathcal{V} = \{0, 1, \dots, n, n+1\}$ and arcs $\mathcal{A} \subseteq \{(i, j) \in \mathcal{V} \times \mathcal{V} : i \neq j, (i, j) \neq (0, n+1)\}$.² Here, nodes 0 and $n+1$ respectively represent the origin and destination depots of vehicle routes but refer to identical physical locations. We use $\mathcal{V}_C = \mathcal{V} \setminus \{0, n+1\}$ to denote the subset of nodes where deliveries must be made. Each node $i \in \mathcal{V}$ is associated with a demand $q_i > 0$ and a service time $s_i \geq 0$. Without loss of generality, we assume that $q_0 = q_{n+1} = 0$ and $s_0 = s_{n+1} = 0$. The travel time along arc $(i, j) \in \mathcal{A}$ and its distance are denoted by $t_{ij} > 0$ and $d_{ij} > 0$, respectively. We assume throughout that the arc lengths satisfy the triangle inequality: $d_{ij} + d_{jk} \geq d_{ik}$ for all $(i, j), (j, k), (i, k) \in \mathcal{A}$.

²Note that \mathcal{G} does not necessarily represent the real-world road network. It is generally constructed by transforming the road network graph into a (possibly complete) graph whose nodes are the vertices of the road network corresponding only to the depot and customers.

We let $\mathcal{K} = \{k_h, k_c\}$ denote the set of vehicle types, where k_h and k_c stand for hybrid REEV and CV types, respectively. There are $m_k \in \mathbb{N}$ vehicles available of type $k \in \mathcal{K}$, each with a carrying capacity of $Q > 0$ demand units and a maximum allowed working time of $T > 0$ units. We assume that all REEVs start from the depot with fully charged batteries and are not allowed to stop en route to recharge them. The per-mile cost of operating a CV is denoted by c_0 , whereas the per-mile costs of operating an REEV in EV and non-EV (i.e., in internal combustion engine or charge-sustaining) modes are denoted by c_e and c_g , respectively. Throughout, we assume that $c_e < c_g \leq c_0$, reflecting the relatively higher costs of fossil fuels (gasoline or CNG) compared with battery costs in EV mode (e.g., electricity for charging). The maximum distance that an REEV can travel in EV mode (also known as its all-electric range) is constrained by $D_E \geq 0$.

We let \mathcal{P} denote the set of all finite and open paths in G that start at the origin depot 0, traverse through a subset of \mathcal{V}_C , and end at the destination depot $n+1$. Specifically, each path $p = (v_0 = 0, v_1, v_2, \dots, v_l, v_{l+1} = n+1) \in \mathcal{P}$ satisfies $l > 0$, $\mathcal{A}_p := \{(v_h, v_{h+1}) : h \in \{0, 1, \dots, l\}\} \subseteq \mathcal{A}$ and $v_i \in \mathcal{V}_C$ for $i \in \{1, \dots, l\}$. We use $\eta_p = \sum_{i=1}^l q_{v_i}$ to denote the total demand along p , $\delta_p = \sum_{i=0}^l d_{v_i v_{i+1}}$ to denote the travel distance along p and $\tau_p = \sum_{i=0}^l (t_{v_i v_{i+1}} + s_{v_i})$ to denote the total duration along p (that is, the total time to process all nodes in p). We use the notation $i \in p$ to indicate that node i belongs to path p ; that is, $i = v_j$ for some $j \in \{0, 1, \dots, l, l+1\}$.

If a path $p \in \mathcal{P}$ is traversed by a vehicle of type $k \in \mathcal{K}$, then the cost incurred is given as follows:

$$c_{pk} = \begin{cases} c_e \min\{\delta_p, D_E\} + c_g \max\{\delta_p - D_E, 0\} & \text{if } k = k_h \\ c_0 \delta_p & \text{if } k = k_c. \end{cases} \quad (1)$$

The first term in the first condition on the right-hand side states that REEVs ($k = k_h$) incur a mileage cost of c_e in EV mode but that the total distance they can travel in this mode is restricted by their EV range, D_E . The second term enforces a mileage cost of c_g in non-EV mode for any travel distance that exceeds D_E . The second condition on the right-hand side of [eq. \(1\)](#) states that CVs ($k = k_c$) always incur a per-mile cost of c_0 .

A path that is elementary and both time- and capacity-feasible is also called a vehicle route (or simply a route), and we use $\mathcal{R} \subseteq \mathcal{P}$ to denote the set of all such routes. Specifically, each route $r = (v_0 = 0, v_1, v_2, \dots, v_l, v_{l+1} = n+1) \in \mathcal{R}$ satisfies the following conditions:

(C1) r is elementary: $v_i \neq v_j$ for $i, j \in \{1, \dots, l\}$, $i \neq j$.

(C2) r satisfies the capacity limit: $\eta_r \leq Q$.

(C3) r satisfies the duration limit: $\tau_r \leq T$.

A vector of F vehicle routes $R = (r_1, \dots, r_F)$ in conjunction with a fleet composition vector $\kappa = (\kappa_1, \dots, \kappa_F)$, where $r_f \in \mathcal{R}$, $\kappa_f \in \mathcal{K}$ for $f \in \{1, \dots, F\}$, is said to define a feasible solution of the REEVRP if the following are satisfied:

(D1) Routes r_1, \dots, r_F partition \mathcal{V}_C ; that is, each customer node is visited on exactly one route.

(D2) The number of routes served by vehicles of type k does not exceed their available number; that is $\sum_{f=1}^F \mathbb{I}[\kappa_f = k] \leq m_k$ for all $k \in \mathcal{K}$, where $\mathbb{I}[\mathcal{E}]$ is the indicator function that evaluates to 1 if expression \mathcal{E} is true and 0 otherwise.

The cost of a feasible solution is the sum of the costs of its individual routes, $c(R, \kappa) := \sum_{f=1}^F c_{r_f \kappa_f}$, and the goal is to determine a feasible solution of minimum cost.

4.2 Set Partitioning Model

The central challenge in solving the REEVRP is to model the nonlinear costs in [eq. \(1\)](#) incurred by REEVs using a tractable linear integer programming formulation. For this purpose we first define an extended set of vehicle types $\mathcal{K}^+ := \{k_e, k_g, k_c\}$. Here the vehicle type k_c refers to CV types as before. However, the hybrid REEV type k_h is represented by using two new types that we denote as k_e and k_g . The subtype k_e represents vehicles that operate purely in EV mode at all times; that is, their total individual travel distances do not

exceed D_E units. The subtype k_g represents vehicles that necessarily switch to non-EV mode at some point; that is, their total individual travel distances exceed D_E units. As we will show, this partition enables a linear cost representation that can be modeled in an integer linear program.

Our set partitioning model uses binary variables λ_{rk} that indicate whether a route r is used in the optimal solution and whether it is performed by a vehicle of type $k \in \mathcal{K}^+$. It is well known (e.g., see [Poggi and Uchoa \[2014\]](#)) that one does not sacrifice optimality if these variables are defined over routes r that are not necessarily elementary, that is, those that may not satisfy condition (C1). In particular, we define these variables over the set $\mathcal{R}^{ng} \subseteq \mathcal{P}$ of feasible *ng*-routes (e.g., see [Baldacci et al. \[2011\]](#)), which are defined as follows. For each $i \in \mathcal{V}_C$, let $NG_i \subseteq \mathcal{V}_C$ denote the *ng*-set of node i . This set may be defined to consist of any subset of \mathcal{V}_C , with the only restriction being that $i \in NG_i$. It is common to define NG_i to consist of the Δ -nearest neighbors of i , with respect to some metric such as the one induced by the distance matrix (usually $\Delta \in \{8, 9, \dots, 16\}$). Then, each route $r = (v_0 = 0, v_1, v_2, \dots, v_l, v_{l+1} = n + 1) \in \mathcal{R}^{ng}$ satisfies conditions (C1'), (C2) and (C3), where the elementarity condition is relaxed as follows:

(C1') r is *ng*-elementary: if $v_i = v_j$ for some $i, j \in \{1, \dots, l\}$, $i < j$, then there exists v_k for some $k \in \{i + 1, i + 2, \dots, j - 1\}$ such that $v_i \notin NG_{v_k}$.

We can now present our IP model. In the following formulation, for each path $p = (v_0 = 0, v_1, v_2, \dots, v_l, v_{l+1} = n + 1) \in \mathcal{P}$, integer-valued parameters α_{ip} and β_{ijp} count the number of times node i and arc (i, j) , respectively, appear on a (not necessarily elementary) path $p \in \mathcal{P}$; whereas $\mathcal{C}_p = \{(v_i, v_j) \in \mathcal{A} : i, j \in \{0, 1, \dots, l, l + 1\}, i < j\}$ denotes the set of arcs in the transitive closure of (the subgraph induced by) p .

$$\text{minimize}_{\lambda} \sum_{r \in \mathcal{R}^{ng}} (c_e \delta_r \lambda_{rk_e} + [c_g \delta_r - (c_g - c_e) D_E] \lambda_{rk_g} + c_0 \delta_r \lambda_{rk_c}) \quad (2a)$$

$$\text{subject to} \sum_{k \in \mathcal{K}^+} \sum_{r \in \mathcal{R}^{ng}} \alpha_{ir} \lambda_{rk} = 1, \quad i \in \mathcal{V}_C \quad (2b)$$

$$\sum_{r \in \mathcal{R}^{ng}} (\lambda_{rk_e} + \lambda_{rk_g}) \leq m_{k_h} \quad (2c)$$

$$\sum_{r \in \mathcal{R}^{ng}} \lambda_{rk_c} \leq m_{k_c} \quad (2d)$$

$$\sum_{(i,j) \in \mathcal{C}_p} \sum_{r \in \mathcal{R}^{ng}} \beta_{ijr} \lambda_{rk_g} \leq |\mathcal{A}_p| - 1, \quad p \in \mathcal{R} : \delta_p \leq D_E \quad (2e)$$

$$\lambda_{rk_e} = 0, \quad r \in \mathcal{R}^{ng} : \delta_r > D_E \quad (2f)$$

$$\lambda_{rk} \in \{0, 1\}, \quad r \in \mathcal{R}^{ng}, \quad k \in \mathcal{K}^+ \quad (2g)$$

In this formulation, [constraint \(2b\)](#) ensures that each customer is visited exactly once, whereas constraints [\(2c\)–\(2d\)](#) ensure that the fleet size is respected. Next, [constraint \(2e\)](#) is an *infeasible path elimination constraint* (IPEC), which enforces that hybrid REEVs of subtype k_g are not assigned to routes whose travel distances are less than or equal to D_E units (a formal proof is provided shortly). Its counterpart [constraint \(2f\)](#) ensures that hybrid vehicles of subtype k_e are not assigned to routes whose travel distances exceed D_E units. [constraint \(2g\)](#) requires that the variables be binary.

The following propositions establish the validity of formulation (2).

Proposition 1. *Every feasible solution of the REEVRP can be mapped to a feasible solution of formulation (2) with the same cost.*

Proof. Let (R, κ) be a feasible solution of the REEVRP where $R = (r_1, \dots, r_F)$ and $\kappa = (\kappa_1, \dots, \kappa_F)$ with $r_f \in \mathcal{R}$ and $\kappa_f \in \mathcal{K}$ for each $f \in \{1, \dots, F\}$. By definition, (R, κ) satisfies conditions (C1), (C2), (C3), (D1), and (D2). Construct λ as follows: For each $f \in \{1, \dots, F\}$, define

$$\lambda_{r_f k_c} = \mathbb{I}[\kappa_f = k_c], \quad \lambda_{r_f k_e} = \mathbb{I}[\kappa_f = k_h] \mathbb{I}[\delta_{r_f} \leq D_E], \quad \lambda_{r_f k_g} = \mathbb{I}[\kappa_f = k_h] \mathbb{I}[\delta_{r_f} > D_E].$$

Also, define $\lambda_{rk} = 0$ for all $r \in \mathcal{R}^{ng} \setminus \{r_1, \dots, r_F\}$ and all $k \in \mathcal{K}^+$. Now observe that λ satisfies [constraint \(2b\)](#) because R satisfies condition (D1) and because $\alpha_{ir_f} = \mathbb{I}[i \in r_f]$ (since each r_f is elementary by

definition (C1)); it also satisfies constraints (2c)–(2d) because R satisfies condition (D2); and similarly it satisfies constraints (2f)–(2g) by construction of λ .

To see that constraint (2e) is also satisfied, fix any route $p = (v_0 = 0, v_1, \dots, v_\ell, v_{\ell+1} = n+1) \in \mathcal{R}$ such that $\delta_p \leq D_E$, and define the flow variables:

$$x_{ij} = \begin{cases} \sum_{r \in \mathcal{R}^{ng}} \beta_{ijr} \lambda_{rk_g} & \text{if } (i, j) \in \mathcal{C}_p \\ 0 & \text{otherwise} \end{cases} \quad \text{for all } (i, j) \in \mathcal{V} \times \mathcal{V}.$$

By construction of λ and because of the elementarity (C1) and degree (D1) conditions, observe that $x_{ij} \in \{0, 1\}$ and $x_{ij} = 1$ if and only if arc (i, j) is part of some route $r \in \mathcal{R}$ such that $\delta_r > D_E$ and it is performed by a vehicle of type k_h . Also, for any $i \in \mathcal{V}_C$, since the definition of α and β implies that

$$\alpha_{ir} = \sum_{j \in \mathcal{V}_C \cup \{n+1\} : (i, j) \in \mathcal{A}} \beta_{ijr} = \sum_{j \in \mathcal{V}_C \cup \{0\} : (j, i) \in \mathcal{A}} \beta_{jir},$$

the degree condition (D1) requires that $\sum_{j \in \mathcal{V}} x_{ij} \leq 1$ and $\sum_{j \in \mathcal{V}} x_{ji} \leq 1$. Along with $x_{ij} \in \{0, 1\}$, this implies that in the subgraph $G_x = \{\cup_{i=0}^{\ell+1} \{v_i\}, \mathcal{A}_x\}$ induced by $\mathcal{A}_x := \{(i, j) \in \mathcal{A} : x_{ij} = 1\}$, each of the nodes v_1, v_2, \dots, v_ℓ has out-degree and in-degree equal to 0 or 1. We now claim that the terminal nodes $u_1, u_2, \dots, u_m \in p$ of each of the $m := \sum_{j=1}^{\ell} x_{0v_j}$ subpaths in G_x starting from node 0 satisfy $u_i \neq n+1$. Indeed, if some $u_i = n+1$, then this implies that there is a path $\pi \in \mathcal{P}$ for which $\lambda_{\pi k_g} = 1$ (by construction of λ and x), and hence $\delta_\pi > D_E$. But since the travel distances satisfy the triangle inequality and π is contained in the transitive closure \mathcal{C}_p of p , we also must have $\delta_\pi \leq \delta_p \leq D_E$, resulting in a contradiction. Therefore, $u_1, u_2, \dots, u_m \in \mathcal{V}_C$, and we have

$$\begin{aligned} \sum_{(i,j) \in \mathcal{C}_p} x_{ij} &= \sum_{j=1}^{\ell} x_{0v_j} + \sum_{i=1}^{\ell} \sum_{j=i+1}^{\ell+1} x_{v_i v_j} = m + \sum_{v \in \{v_1, \dots, v_\ell\} \setminus \{u_1, \dots, u_m\}} \sum_{j: (v, j) \in \mathcal{C}_p} x_{vj} \\ &\leq m + \sum_{v \in \{v_1, \dots, v_\ell\} \setminus \{u_1, \dots, u_m\}} \sum_{j \in \mathcal{V}} x_{vj} \\ &\leq m + (\ell - m) \left(\text{since } \sum_{j \in \mathcal{V}} x_{ij} \leq 1 \right) \\ &= |\mathcal{A}_p| - 1, \end{aligned}$$

showing that constraint (2e) is also satisfied.

The objective value of the constructed λ solution is equal to $\sum_{f=1}^F \text{cost}_f$, where

$$\begin{aligned} \text{cost}_f &= \begin{cases} c_e \delta_{r_f} \mathbb{I}[\delta_{r_f} \leq D_E] + [c_g \delta_{r_f} - (c_g - c_e) D_E] \mathbb{I}[\delta_{r_f} > D_E] & \text{if } \kappa_f = k_h \\ c_0 \delta_{r_f} & \text{if } \kappa_f = k_c \end{cases} \\ &= \begin{cases} c_e \min\{\delta_{r_f}, D_E\} + c_g \max\{\delta_{r_f} - D_E, 0\} & \text{if } \kappa_f = k_h \\ c_0 \delta_{r_f} & \text{if } \kappa_f = k_c \end{cases} \\ &= c_{r_f \kappa_f}, \end{aligned}$$

where the first equality follows by construction of λ and the second follows by directly comparing both sides of the equality separately in the cases $\delta_{r_f} \leq D_E$ and $\delta_{r_f} > D_E$. \square

Proposition 2. *Every feasible solution of formulation (2) can be mapped to a feasible solution of the REEVRP with the same cost.*

Proof. Suppose that λ is a feasible solution of formulation (2). Construct (R, κ) as follows: (i) $F \leftarrow 0$; (ii) for all $r \in \mathcal{R}^{ng}$, if $\sum_{k \in \mathcal{K}^+} \lambda_{rk} > 0$, then set $F \leftarrow F + 1$, $r_f \leftarrow r$, and $\kappa_f \leftarrow k_c$ if $\lambda_{rk_c} = 1$ and $\kappa_f \leftarrow k_h$ otherwise. We claim that (R, κ) is a feasible solution of the REEVRP with the same cost as the objective

value of λ . First, observe that all individual routes r_f satisfy conditions (C2) and (C3) by definition of their domain \mathcal{R}^{ng} . Second, the degree constraint (2b) implies that these routes also satisfy condition (C1), and hence the solution (R, κ) satisfies condition (D1). Third, the fleet size constraints (2c)–(2d) imply that condition (D2) is also satisfied by (R, κ) . To see that its cost is the same as the objective value of λ , consider any route r_f . We consider three possibilities. If $\kappa_f = k_c$, then $\lambda_{r_f k_c} = 1$ (by construction), and its objective contribution is $c_0 \delta_{r_f} = c_{r_f \kappa_f}$ (see eq. (1)). If $\kappa_f = k_h$ and $\delta_{r_f} \leq D_E$, then the IPEC (2e) for $p = r_f$ ensures that $\lambda_{r_f k_g} = 0$ (hence $\lambda_{r_f k_e} = 1$), implying that its objective contribution is $c_e \delta_{r_f}$, which is the same as $c_{r_f \kappa_f}$ (see eq. (1)). If $\kappa_f = k_h$ but $\delta_{r_f} > D_E$, then constraint (2f) ensures that $\lambda_{r_f k_e} = 0$ (hence $\lambda_{r_f k_g} = 1$) and its corresponding objective contribution is $c_g \delta_{r_f} - (c_g - c_e) D_E = c_e \min\{\delta_{r_f}, D_E\} + c_g \max\{\delta_{r_f} - D_E, 0\} = c_{r_f \kappa_f}$ (see eq. (1)). This proves the claim. \square

4.3 Branch-Price-and-Cut Algorithm

The number of binary variables and constraints in formulation (2) scales exponentially with n , prohibiting its direct solution by conventional IP solvers. To address this issue, one can solve formulation (2) using a BPC algorithm, which embeds column generation (also called pricing) and cutting planes in each node of a branch-and-bound search tree. In the following paragraphs we describe the details of such an algorithm as it applies to formulation (2); see Pecin et al. [2017a], Pessoa et al. [2020] for general references on BPC algorithms for vehicle routing.

4.3.1 Pricing subproblems

The linear programming (LP) relaxations of (2) at each node of the branch-and-bound tree are solved by column generation. The latter starts with only a small subset of feasible ng -routes for each vehicle type $k \in \mathcal{K}^+$ and then iteratively solves pricing subproblems to add back those routes that have been neglected but that improve the objective value of the current LP relaxation. For each vehicle type $k \in \mathcal{K}^+$, the pricing subproblem is a shortest-path problem with resource constraints (SPPRCs) where the objective is to identify feasible ng -routes with negative reduced costs [Pugliese and Guerriero, 2013]. Each SPPRC is defined on the original graph G where arc $(i, j) \in \mathcal{A}$ has a travel cost equal to $c_e d_{ij}$ for vehicles of type k_e , $c_0 d_{ij}$ for vehicles of type k_c , and $c_g d_{ij} - \mathbb{I}[i = 0](c_g - c_e) D_E$ for vehicles of type k_g , each of which is appropriately modified to account for contributions coming from the dual values of constraints in the current LP relaxation. In addition, we define the following resource constraints for the various vehicle types:

- Capacity resource constraint for vehicles of type $k \in \mathcal{K}^+$, for which arc $(i, j) \in \mathcal{A}$ has a resource consumption of q_j units and each node $i \in \mathcal{V}$ has resource limits of $[0, Q]$
- Duration resource constraint for vehicles of type $k \in \mathcal{K}^+$, for which arc $(i, j) \in \mathcal{A}$ has a resource consumption of $(t_{ij} + s_i)$ units and each node $i \in \mathcal{V}$ has resource limits of $[0, T]$
- Travel distance resource constraint for vehicles of type k_e , for which arc $(i, j) \in \mathcal{A}$ has a resource consumption of d_{ij} units and each node $i \in \mathcal{V}$ has resource limits of $[0, D_E]$

Although the SPPRC is weakly NP-hard, it can be solved efficiently in pseudo-polynomial time by using a labeling-based dynamic programming algorithm [Pugliese and Guerriero, 2013], the tractability of which depends crucially on the efficiency of its dominance checks. In this context, all of the above resources are “disposable,” which enables the implementation of efficient dominance checks [Pessoa et al., 2020, Section 5.1.2]. Note that a partial path p (“partial” in the sense that it does not necessarily end at node $n + 1$) dominates another partial path p' whenever for every feasible path extended from p' we can find a corresponding feasible path extension from p with a smaller reduced cost. More formally, p dominates p' if both end at the same terminal node and satisfy $F_p \subseteq F_{p'}$, $\eta_p \leq \eta_{p'}$ and $\tau_p \leq \tau_{p'}$ (and additionally $\delta_p \leq \delta_{p'}$ for vehicles of type k_e), where F_p is the set of so-called forbidden vertices of path p [Baldacci et al., 2011].

In principle, the *lower travel distance limit* of D_E units on vehicles of subtype k_g can also be handled by defining an additional resource whose consumption along arc $(i, j) \in \mathcal{A}$ is d_{ij} and whose resource limits are equal to $[0, \infty]$ for nodes $i \in \{0\} \cup \mathcal{V}_C$ and equal to $(D_E, \infty]$ for the destination node $n + 1$. Doing so, however, amounts to defining a “nondisposable” resource, since an accumulated resource consumption of less than D_E units at node $n + 1$ is not feasible. In this case the corresponding dominance check is weak:

a partial path p dominates p' only if $\delta_p = \delta_{p'}$ (along with the aforementioned conditions of ending at the same terminal node, $F_p \subseteq F_{p'}$, $\eta_p \leq \eta_{p'}$ and $\tau_p \leq \tau_{p'}$) [Pessoa et al., 2020, Section 5.1.2]. This reduces the SPPRC to almost complete enumeration of (exponentially many) paths with distances greater than D_E units, resulting in impractical memory requirements. Therefore, the lower travel distance limit for vehicles of subtype k_g is enforced by using cutting planes via the IPEC defined in [constraint \(2e\)](#), which we describe in [section 4.3.2](#).

Because dominance checks can become computationally expensive, labels (or partial paths) with similar resource consumption levels are usually saved in a so-called bucket, and dominance is checked only among labels in the same bucket [Pecin et al., 2017a, Sadykov et al., 2021]. Other acceleration techniques including heuristic pricing [Fukasawa et al., 2006], bidirectional labeling [Righini and Salani, 2008], variable fixing [Irnich et al., 2010], stabilized column generation [Pessoa et al., 2018], and dynamic ng -sets [Roberti and Mingozzi, 2014] have also been proposed.

4.3.2 Cutting planes and separation

The solution of a restricted LP at each node of the search tree is iteratively followed by cut separation, where the goal is to add valid cutting planes that are violated at the current LP solution. These include cuts that are necessary to ensure solution feasibility, such as the IPEC [constraint \(2e\)](#), as well as those that are only strengthening but not necessary.

Although the family of the IPEC has exponentially many inequalities, we can separate them using the following polynomial-time enumerative path-growing procedure. Given a (fractional) solution λ^* of formulation (2), we first construct the corresponding graph $G_x = (\mathcal{V}, \mathcal{A}_x)$, where $\mathcal{A}_x = \{(i, j) \in \mathcal{A} : x_{ij}^* > 0\}$, and the classical flow variables x^* are defined as follows:

$$x_{ij}^* = \begin{cases} \sum_{r \in \mathcal{R}^{ng}} \beta_{ijr} \lambda_{rk_g}^* & \text{if } (i, j) \in \mathcal{A} \\ 0 & \text{otherwise} \end{cases} \quad \text{for all } (i, j) \in \mathcal{V} \times \mathcal{V}.$$

Starting from node 0, elementary paths are grown in a depth-first fashion by moving along the incident outgoing arcs of strictly positive flow. Each path $p = (v_0 = 0, v_1, v_2, \dots, v_l, v_{l+1})$ is extended as long as the total flow on its transitive closure, $\sum_{i=0}^l \sum_{j=i+1}^{l+1} x_{v_i v_j}^*$, is strictly greater than $|\mathcal{A}_p| - 1 = l$ and as long as $v_{l+1} \neq n + 1$; that is, the path has not reached the destination depot. By construction, every such identified path p satisfying $\delta_p \leq D_E$ defines a corresponding violated IPEC [constraint \(2e\)](#).

Apart from the above necessary IPEC, we also separate the (strengthening but not necessary) rounded capacity inequalities:

$$\sum_{(i,j) \in \mathcal{A} \cap (S \times S)} \sum_{k \in \mathcal{K}^+} \sum_{r \in \mathcal{R}^{ng}} \beta_{ijr} \lambda_{rk} \leq |S| - \left\lceil \frac{1}{Q} \sum_{i \in S} q_i \right\rceil, \quad S \subseteq \mathcal{V}_C. \quad (3)$$

Furthermore, we add the following strengthening inequality to the initial master problem:

$$\sum_{r \in \mathcal{R}^{ng}} \delta_r \lambda_{rk_g} \geq D_E \sum_{r \in \mathcal{R}^{ng}} \lambda_{rk_g}. \quad (4)$$

[Equation \(4\)](#) states that the total distance covered by all hybrid vehicles of subtype k_g (i.e., the left-hand side) must be at least as large as D_E times the total number of those vehicles used (i.e., the right-hand side). We emphasize that adding this inequality along with the infeasible path elimination and rounded capacity constraints does not affect the complexity of the pricing subproblems, since their corresponding dual values affect only the arc costs defining the SPPRC; that is, they are “robust cuts.” In contrast, the addition of so-called limited-memory subset row or rank-1 cuts (which are also strengthening but not necessary) significantly increases the complexity of the pricing subproblems; see Jepsen et al. [2008] and Pecin et al. [2017a,b] for details.

4.3.3 Other implementation details

Pricing and cut generation are iteratively performed at every node of the search tree until either the resulting (restricted) LP solution is integral and feasible or the addition of cuts does not significantly improve the LP objective value (i.e., tailing off occurs). At this point, the algorithm proceeds by route enumeration or by branching. The former refers to enumerating all elementary routes with reduced costs less than the current (local) primal-dual optimality gap, whenever this gap is sufficiently small (note that a dual bound can be obtained from any exact solution of the pricing subproblem). The complexity and implementation of this enumeration step are similar to that of the aforementioned labeling algorithm [Baldacci et al., 2008]. If the optimality gap is large, then we proceed to branching, where we prioritize branching on the assignment of customers to vehicle types, given by $\sum_{r \in \mathcal{R}^{ng}} \alpha_{ir} \lambda_{rk}$. That is, we branch primarily on the left-hand side of [constraint \(2b\)](#) for a fixed $k \in \mathcal{K}^+$ and secondarily on the standard flow variable $\sum_{k \in \mathcal{K}^+} \sum_{r \in \mathcal{R}^{ng}} \beta_{ijr} \lambda_{rk}$.

We utilize the VRPSolver package [Pessoa et al., 2020] in our implementation with default options except for the addition of at most 10 limited-arc-memory rank-1 cuts (per cut round), bidirectional labeling for both heuristic and exact pricing, cut tailing-off threshold of 1%, and soft and hard time thresholds of 5 and 10 seconds, respectively, in the labeling algorithm.

5 Iterated Tabu Search Metaheuristic

This section describes an iterated tabu search metaheuristic tailored to solving large-scale REEVRP instances (e.g., with more than a thousand nodes). This approach is inspired by existing local-search-based metaheuristics for the heterogeneous VRP [Penna et al., 2013, Subramanyam et al., 2020], albeit with some important differences that exploit characteristics of the REEVRP. A key feature of this ITS approach is its simplicity because it introduces few user-defined parameters and does not require any instance-specific features to accelerate the search process.

ITS can be viewed as an improved version of iterated local search [Lourenço et al., 2010]. The latter refers to the repeated application of local search to a current solution that may be generated either from scratch using a construction heuristic or from perturbing an existing solution. Local search refers to the repeated use of fundamental moves that transform the current solution into a neighbor solution [Aarts and Lenstra, 2003]. Given a set M of moves, let $\Omega_M(R, \kappa)$ be the neighborhood of the current solution (R, κ) with respect to the move set M ; that is, $\Omega_M(R, \kappa)$ is the set of solutions that can be reached from (R, κ) by applying a move from M . The crucial components of any local-search-based metaheuristic, therefore, are the definition of M and the exploration of $\Omega_M(R, \kappa)$ using a search algorithm. For the latter, we use tabu search [Glover, 1997] since it provides an efficient mechanism to escape local minima in the neighborhood, $\Omega_M(R, \kappa)$. For the former, we consider M to be the (intra- and interroute) 1-0 relocate, 1-1 exchange, and 2-opt moves that involve the deletion and reinsertion of nodes or edges. These moves define corresponding node- and edge-exchange neighborhoods that belong to the family of k -Opt and λ -interchange neighborhoods [Funke et al., 2005] with sizes $|\Omega_M(R, k)| = O(n^2)$.

[Algorithm 1](#) presents the ITS metaheuristic. It consists of a construction phase and a perturbation phase. The former involves the construction of an initial solution that is then improved by using tabu search. The latter involves iteratively perturbing this solution and again improving it by using tabu search. If this phase fails to encounter a solution that is better than the best solution found in the current iteration, (R', κ') , for more than θ^{restart} attempts, then the construction phase is restarted. The overall algorithm is terminated when the total run time or the number of iterations exceeds some prescribed limit.

The merit function $\phi(R, \kappa)$ is used to compare different solutions. When $R = (r_1, \dots, r_F)$ and $\kappa = (\kappa_1, \dots, \kappa_F)$, we define this function as the weighted sum of its cost and constraint violations:

$$\phi(R, \kappa) = c(R, \kappa) + \sum_{f=1}^F (\varphi^Q \max\{\eta_{r_f} - Q, 0\} + \varphi^T \max\{\tau_{r_f} - T, 0\}), \quad (5)$$

where φ^Q and φ^T are penalty parameters that are chosen to be sufficiently large to ensure that a solution (R, κ) is better than another one (R', κ') (i.e., $\phi(R, \kappa) < \phi(R', \kappa')$) only if it primarily has a lower or equal constraint violation and secondarily a lower cost. This approach allows the algorithm to handle infeasible solutions whose routes may not satisfy the capacity [\(C2\)](#) or duration [\(C3\)](#) conditions. Moreover, large values

Algorithm 1: Iterated tabu search metaheuristic.

Input: User-defined parameters θ^{rcl} , θ^{tenure} , θ^{maxiter} , θ^{restart} , θ^{perturb}
Output: Best found solution (R^B, κ^B)

```
1  $(R^B, \kappa^B) \leftarrow (\emptyset, \emptyset)$ 
2 while termination criteria not met do
3    $(R, \kappa) \leftarrow \text{CONSTRUCT SOLUTION}(\theta^{\text{rcl}})$  /* Construction phase */
4    $(R, \kappa) \leftarrow \text{TABU SEARCH}((R, \kappa), \theta^{\text{tenure}}, \theta^{\text{maxiter}})$ 
5   counter  $\leftarrow 0$ ,  $(R', \kappa') \leftarrow (R, \kappa)$  /* Perturbation phase */
6   while counter  $< \theta^{\text{restart}}$  do
7      $(R, \kappa) \leftarrow \text{PERTURB SOLUTION}((R, \kappa), \theta^{\text{perturb}})$ 
8      $(R, \kappa) \leftarrow \text{TABU SEARCH}((R, \kappa), \theta^{\text{tenure}}, \theta^{\text{maxiter}})$ 
9     if  $\phi(R, \kappa) < \phi(R^B, \kappa^B)$  then  $(R^B, \kappa^B) \leftarrow (R, \kappa)$ 
10    if  $\phi(R, \kappa) < \phi(R', \kappa')$  then
11      counter  $\leftarrow 0$ 
12       $(R', \kappa') \leftarrow (R, \kappa)$ 
13    else
14      counter  $\leftarrow \text{counter} + 1$ 
```

of φ^Q and φ^T discourage the algorithm from re-entering the infeasible region as soon as a feasible solution is encountered. The parameters θ^{rcl} , θ^{tenure} , θ^{maxiter} , and θ^{perturb} are used as inputs to the construction heuristic, tabu search, and perturbation mechanisms, respectively, which we describe next.

5.1 Construction Heuristic

The CONSTRUCT SOLUTION procedure constructs an REEVRP solution that always satisfies the set partitioning (D1) and fleet availability (D2) constraints; however, the individual routes need not necessarily satisfy the vehicle capacity (C2) or duration (C3) limits. The procedure adds customers into an empty solution using the following sequential insertion heuristic.

Each iteration of the insertion heuristic consists of two steps. First, if the number of REEV routes in the current (partial) solution is less than its available number m_{k_h} , then an empty REEV route is constructed; otherwise, an empty CV route is constructed provided again that the number of CV routes is less than m_{k_c} . Second, unrouted customers are repeatedly inserted into this (current) route until either all customers have been routed or there is no feasible insertion position in the route; in other words, inserting any unrouted customer at any position in the route would exceed either the vehicle capacity Q or the duration limit T . Specifically, each unrouted customer that can be feasibly inserted into at least one position is first added to a so-called restricted candidate list. A random customer is then selected from this list and inserted in the position that minimizes a greedy function, defined as a randomly weighted sum of the corresponding insertion cost, the residual capacity $Q - \eta_r$, and the residual duration $T - \tau_r$. The restricted candidate list is cardinality-based; that is, it is sorted according to the greedy function, and its length is not allowed to exceed θ^{rcl} . This parameter controls the extent of randomization and greediness with our implementation using $\theta^{\text{rcl}} = 3$. If no empty route can be constructed in the first step, then any remaining unrouted customer is inserted into an existing route (and corresponding position) for which a randomly weighted sum of the insertion cost, capacity violation, and duration violation is minimized.

5.2 Tabu Search

The TABU SEARCH procedure enhances the performance of classical local search by (i) allowing non-improving moves and (ii) potentially disallowing improving moves. This enhancement, which is achieved by using a short-term memory that is also known as a tabu list, provides a formal mechanism for the current solution to potentially escape from local minima in the neighborhoods defined with respect to the considered elementary moves. The tabu list keeps track of the most recently visited solutions in the search history and prevents revisiting them for a predefined number of local search iterations, θ^{tenure} . Any potential solution that has

been visited in the last θ^{tenure} iterations is marked “tabu” (or forbidden) and inserted into the tabu list, to prevent the algorithm from cycling and repeatedly visiting the same solutions. In fact, a neighbor solution that is in the tabu list can be visited only if certain *aspiration criteria* are met; in particular, the tabu status of a solution is overridden only if it improves upon the best-encountered solution. The tabu search terminates if it performs θ^{maxiter} local search iterations without observing any further improvement. Our implementation uses $\theta^{\text{tenure}} = 20$ and $\theta^{\text{maxiter}} = 100$.

At each iteration, (i) we randomly select a neighborhood (intra- or interroute relocate, exchange, or 2-opt); (ii) traverse it in lexicographic order applying pruning mechanisms based on both feasibility and gain; and (iii) use a first-improvement strategy to replace the current solution. The per-iteration run time is $\mathcal{O}(n^2)$ in the worst case, since the size of each of the aforementioned neighborhoods is $\mathcal{O}(n^2)$ and each move is evaluated in constant time by simply updating the current value of $\phi(R, \kappa)$. Specifically, in addition to storing the values of δ_{r_f} , η_{r_f} and τ_{r_f} for the routes r_f , $f \in \{1, \dots, F\}$ of the current solution, we store and incrementally update (in constant time) the value of $\max\{\delta_{r_f} - D_E, 0\}$ for REEV routes r_f with $\kappa_f = k_h$, which enables the efficient evaluation of the travel cost for REEV routes as well (see eq. (1)).

5.3 Perturbation Mechanism

The goal of the PERTURB SOLUTION procedure is to apply “large” perturbations to the current solution (R, κ) so as to encourage a low likelihood of encountering the resulting solution by applications of tabu search alone. To that end, the procedure first removes the route r of vehicle type k from the current solution that has the maximum average cost per unit of carried load, defined as c_{rk}/η_r . In addition, it removes any routes r'_1, r'_2, \dots that are “sufficiently close” to r , namely, routes for which $\Delta(r, r'_i) := \max_{(i,j) \in r \times r'_i} d_{ij} < \theta^{\text{perturb}}$. If all routes $r' \neq r$ satisfy $\Delta(r, r') \geq \theta^{\text{perturb}}$, then a route $r'' \in \arg \min_{r' \neq r} \Delta(r, r')$ is removed from the current solution. All customers that are visited on the deleted routes are then considered to be unrouted and are added back to the current partial solution by using the same sequential insertion heuristic described in Section 5.1. The parameter θ^{perturb} determines the extent of perturbation, with larger values corresponding to higher extents of perturbation. Our implementation uses $\theta^{\text{perturb}} = 0.6 \max_{(i,j) \in \mathcal{V}_C \times \mathcal{V}_C} d_{ij}$.

6 Case Studies

The goals of this section are threefold. First, Section 6.1 describes the details of our experimental design including data generation. Second, Section 6.2 illustrates the numerical performance of the proposed BPC and ITS algorithms. Third, Sections 6.3 to 6.7 provide key managerial and policy insights.

6.1 Design of experiments

As mentioned in Section 3, we assume that Amazon, FedEx, UPS, and USPS are tasked with delivering packages on a typical weekday in the Chicago metropolitan area. The demand and road network data are obtained from POLARIS [Auld et al., 2016]. In the POLARIS Chicago demand database, 4 million households correspond to 10 million individuals. According to the National Household Travel Survey data from 2017 [Federal Highway Administration, 2017], each household generates roughly one order per week. We therefore assume that on the considered day, roughly 600,000 households (which we also refer to as customers) require an e-commerce delivery service. We randomly draw their locations from the database in a uniform manner.

In addition, we identify 53 depot locations of the aforementioned LSPs from publicly available sources. These sources indicate that Amazon, FedEx, UPS, and USPS operate 17, 10, 13, and 13 depots in the area, respectively. We assume customers are randomly distributed to Amazon, FedEx, UPS, and USPS following 21, 16, 24, and 39 percentage shares, respectively [Spadafora and Rodriguez, 2021]. To efficiently carry out these delivery tasks, we assume that each provider first solves an assignment problem to determine the optimal depot-to-customer allocation. The objective of this assignment problem is to minimize the sum of out-and-back travel times between customers and depots, subject to aggregate depot-level capacity constraints while ensuring that each customer is assigned to some depot.

Each of the 53 depots defines an REEVRP instance averaging 11,447 customers. Figure 2 visualizes an example instance over the Chicago map. These instances are so large that simply constructing and

loading each travel time matrix in memory—let alone optimizing over them—requires more than 130 million shortest-path calculations. This necessitates some sort of network aggregation. Moreover, and perhaps more important, we do not have access to accurate network information at a fine microscopic level. Indeed, since POLARIS is a mesoscopic traffic simulation software, it does not account for local street-level roads where customers are typically located; instead it considers only interstates, principal and other arterials, and major collectors. Therefore, we downscale the original network by aggregating customers located along individual network arcs by introducing nodes known as *superlocations* at the centers of these arcs. Figure 3 details the placement of superlocations on a small portion of the network. After we perform this aggregation, the resulting depot-level instances vary in size from 92 to 3,706 superlocations, with an average and standard deviation of 1,345 and 872, respectively. The travel time matrix between superlocations is computed by using Dijkstra’s shortest paths algorithm, using POLARIS-provided network information (i.e., nodes, arcs, arc-speeds, and arc-lengths).




Figure 2: Layout of an illustrative depot-level REEVRP instance.

The total number of packages to be delivered in the i^{th} superlocation (summed over all customers located there) defines its corresponding demand q_i (since every customer is associated with only one package). The customer-to-customer travel times for customers that are located inside the same superlocation are estimated by dividing their Manhattan distance with a preset low speed of 15 mph because trucks move significantly slower on these minor roads. The optimal visiting sequence of customers in each superlocation is found by solving a small traveling salesman problem. In particular, the sum of the corresponding total travel time and a per-customer package dropping time of σ minutes defines the service time s_i for the i^{th} superlocation. The sum of the corresponding total travel distance is incorporated into the distance matrix entry, d_{ji} , for all other superlocations $j \neq i$.

For each depot, we initially consider only two vehicle types: REEVs and CVs. Subsequently, we also investigate the effects of replacing all REEV trucks with pure BEV trucks, to study the trade-off between these two powertrains. Note that pure BEV trucks can be easily modeled as a special case of REEV trucks by disallowing the use of k_g subtypes in the BPC algorithm and by setting $c_g = +\infty$ (or a sufficiently large number) in the ITS algorithm. Doing so prohibits the hybrid vehicles from using their range extenders. In our experiments we also disallow BEV battery recharging along their routes but increase their EV range to enable a fairer comparison with their REEV counterparts. Nonetheless, since pure BEVs have higher ownership- and infrastructure-related fixed costs that are not accounted for in this study, their estimated costs may be slightly undercounted.




Figure 3: Illustration of superlocations and their relation to other network components.

The mileage cost of diesel-powered CVs is estimated to be $c_0 = c_D/\mu_D$, where c_D is the cost of diesel per gallon (USD/gallon) and μ_D is the fuel economy in miles per gallon (MPG). The mileage cost of REEVs in EV mode is estimated to be $c_e = c_E \cdot \mu_E$, where c_E is the cost of electricity (USD/kWh) and μ_E is the vehicles' electricity consumption rate in EV mode (kWh/mile). Their mileage cost in CNG mode is estimated to be $c_g = c_G/\mu_G$, where c_G is the cost of CNG per gasoline gallon equivalent (USD/GGE) and μ_G is its charge-sustaining mode fuel economy in miles per gasoline gallon equivalent (MPGe). We use the CNG range extender medium-duty vehicle developed by [U.S.-China Clean Energy Research Center \[2021\]](#) as the reference REEV type, for which $\mu_E = 1.14$ kWh/mile, $\mu_G = 10.1$ MPGe, and $D_E = 33$ miles (in the REEV/CV case) and 150 miles (in the BEV/CV case). Based on year 2020 data from [U.S. Energy Information Administration \[2021\]](#) and July 2021 data from [Alternative Fuels Data Center \[2021\]](#), we set $c_E = 9.90$ ¢/kWh, $c_G = 2.22$ USD/GGE, and $c_D = 3.36$ USD/gallon. The average fuel economy of diesel-powered parcel delivery trucks is reported between 7.9 and 10.2 MPG [[Lammert, 2009](#), [Barnitt, 2011](#)]; therefore, we set $\eta_D = 9.0$ MPG. The baseline deployment scenario sets $\sigma = 4$ minutes, $Q = 120$ packages, $T = 10$ hours, and $m_{kh} = 25$ hybrid vehicles per depot (in both REEV/CV and BEV/CV case studies). We investigate regionwide operational impacts and sensitivity of our findings to variations in several of these parameter values in [Sections 6.3 to 6.7](#). In each case study reported, we run the ITS metaheuristic 10 times (each with a time limit of one hour) and use the best of these 10 runs to estimate the corresponding optimal solution.

6.2 Computational Performance of the BPC and ITS Methods

We begin by analyzing the computational performance of the proposed BPC and ITS algorithms as a function of several operational parameters. To that end, we design a testbed of smaller REEVRP instances as follows. First, for each depot in a random subset of 25 depots from [Section 6.1](#), we define two smaller subnetworks by randomly drawing $n \in \{50, 100\}$ superlocations from these instances. We then consider six scenarios, parameterized as per [Table 2](#), for each of the 25×2 subnetworks, resulting in a total of $25 \times 2 \times 6 = 300$ test instances.

All runs were conducted on Intel Xeon servers. The ITS algorithm was implemented in C++, and each ITS run was performed on a 2.3 GHz server, limited to 10 minutes of run time on a single CPU thread. Each BPC run was performed on a 3.1 GHz server, limited to a single CPU thread with 10 GB RAM and 1 hour of computational time. We used the Julia interface of VRPSolver [[Pessoa et al., 2020](#)] as the BPC backend (e.g., for column generation and branching) along with a Julia routine to add cuts. All subordinate optimization problems were solved by using the CPLEX 12.10.0 optimizer [[IBM, 2021](#)]. Each BPC run was

Table 2: Parameterization of operational inputs in scenarios defining the smaller-sized test instances.

Scenario	m_{k_h} (# per depot)	Q (packages)	D_E (miles)	T (hours)
1 (baseline)	5	80	33	8
2 (all REEV)	n	80	33	8
3 (all CV)	0	80	33	8
4 (higher capacity)	5	120	33	8
5 (higher EV range)	5	80	66	8
6 (higher work hours)	5	80	33	10

initialized with an upper bound corresponding to the objective value of a solution that was computed by using a single run of the ITS algorithm.

Table 3 reports the numerical performance of the BPC algorithm across all 300 test instances. For each scenario, the “Optimal” column reports the number of instances that were solved to optimality along with the average solution time in seconds. The “Nonzero optimality gap” column reports the number of instances that could not be solved to optimality in the time limit of one hour, along with the corresponding average optimality gap, defined as $(z_{ub} - z_{lb})/z_{ub} \times 100\%$, where z_{ub} and z_{lb} are the global upper and lower bounds at termination, respectively. These quantities are also graphed in the box plots in Figure 4. The “# IPEC” column reports the average number of violated IPEC constraint (2e) that were added in the algorithm. We observe that the overall performance of the BPC algorithm when compared with the baseline scenario is fairly consistent across scenarios but deteriorates as the EV range D_E or the number of nodes n increases. Nevertheless, we observe that the BPC algorithm can find and certify optimal solutions in 113 out of 300 instances in less than 10 minutes (on average) and can compute reasonably tight optimality gaps of roughly 5% in other cases.

Table 3: Summary of computational performance of the BPC algorithm.

Scenario	# Total Inst.	Optimal		Nonzero gap		# IPEC
		# Inst.	Time (sec)	# Inst.	Gap (%)	
1	50	24	220.6	26	3.20	5.9
2	50	5	572.7	45	10.42	2,175.7
3	50	30	687.1	20	1.57	0.0
4	50	22	499.6	28	2.95	26.1
5	50	10	611.9	40	5.15	71.6
6	50	22	218.2	28	2.73	7.7
$n = 50$	150	102	320.0	48	7.47	602.0
$n = 100$	150	11	1,640.0	139	4.24	160.4
All	300	113	448.5	187	5.07	381.2

Table 4 reports the computational performance of the ITS algorithm across all 300 test instances, with each run repeated 10 times using different random seeds. For each scenario, the “# Optimal” column reports the number of instances for which the optimal objective value is known (from Table 3) and for which the ITS algorithm found a solution matching the optimal objective value. The “Best 10” and “Avg. 10” columns report, respectively, the best and average optimality gaps of the ITS solutions with respect to the global lower bounds from Table 3. These quantities are defined as $(z_{best} - z_{lb})/z_{best} \times 100\%$ and $(z_{avg} - z_{lb})/z_{avg} \times 100\%$, where $z_{best} = \min\{z_1, \dots, z_{10}\}$ and $z_{avg} = (z_1 + \dots + z_{10})/10$ are the best and average of the objective values, z_1, \dots, z_{10} , from the 10 ITS runs. The “Avg. dev. (%)” and “Avg. time (sec)” columns report the average deviation, $\frac{1}{10} \sum_{k=1}^{10} (z_k - z_{best})/z_{best} \times 100\%$ and average time to find the best solution, respectively. These quantities are also graphed in the box plots in Figure 5.

We make the following observations from Table 4. First, the ITS algorithm can find the optimal solution in 91 out of 113 instances for which the optimal value is known. This indicates that the reported optimality gaps in Table 3 and Table 4 are potentially quite conservative and limited by the lower bounds from the BPC algorithm rather than the (suboptimality of the solutions computed by the ITS algorithm. Second, the low deviation of roughly 0.23% across the 10 runs indicates that the ITS method is fairly robust to the




Figure 4: Guaranteed optimality gaps (for instances not solved to optimality in one hour) and solution times (for optimally solved instances) of the BPC algorithm across the 6 scenarios in Table 2.

Table 4: Summary of computational performance of the ITS algorithm.

Scenario	# Total Inst.	# Optimal	Optimality gap		Avg. dev. (%)	Avg. time (sec)
			Best 10	Avg. 10		
1	50	19/24	1.43	1.73	0.32	199.0
2	50	5/5	9.28	9.39	0.12	165.1
3	50	21/30	0.61	0.73	0.11	179.6
4	50	19/22	1.47	1.73	0.26	191.5
5	50	8/10	3.87	4.15	0.29	226.4
6	50	19/22	1.30	1.56	0.26	187.6
$n = 50$	150	90/102	2.41	2.45	0.04	91.8
$n = 100$	150	1/11	3.58	3.98	0.41	291.3
All	300	91/113	2.99	3.22	0.23	191.5




Figure 5: Average deviation and run times of the ITS metaheuristic across the 6 scenarios in Table 2.

initial random seed. Third, the run times are reasonably small across the different scenarios, with the biggest distinction being the problem size, where the run time can increase from roughly 90 seconds for $n = 50$ nodes to 3 minutes for $n = 100$ nodes.

6.3 Impact of EV Deployment on Systemwide Performance Metrics

We now analyze different deployment scenarios to assess the impact of REEVs and BEVs on the operational energy cost, VMT, and VHT. The first case considers all regional-level deliveries (i.e., the 53-depot Chicago metropolitan area) using REEVs and CVs. The second case considers the impact of replacing all REEVs with BEVs. In the following sections we refer to these two cases simply as “REEV” and “BEV,” respectively. In each of the two cases, we consider six scenarios of 0, 25, 50, 100, 200, and 2,000 EVs per depot (be it REEV or BEV). Note that the number 2,000 here is meant to represent an arbitrarily large value; no depot uses so many vehicles. We refer to the scenarios of 0 EVs and 2,000 EVs as “All CV” and “All EV,” respectively.




Figure 6: Impact of EV deployment rate on cost, VMT, VHT, and fleet composition.

Figure 6 summarizes the aggregate regional-level results. First, Figure 6a shows that the cost decrease relative to the All CV scenario, as a function of the number of EVs per depot. We observe that using either REEVs or BEVs can substantially reduce operational costs by up to 67%. Perhaps more important, we see a difference of up to 11% between the REEV and BEV cases in terms of cost reduction. Since BEVs have higher infrastructure- and ownership-related fixed costs, REEVs may be preferable whenever the difference in fixed costs is higher than 11%. Figure 6b and Figure 6c report the regional-level VMT and VHT, respectively. We observe that they exhibit a maximum when the number of EVs is roughly between 100 and 200, depending on the depot. This is not unexpected since the All CV and All EV cases must travel roughly the same number of miles as a homogeneous fleet of vehicles and intermediate compositions must trade off between more EVs of cheaper travel cost and fewer CVs of higher costs. This is supported by Figure 6d and Figure 6e, which break down the VMT per drive mode in REEV and BEV cases, respectively. We observe that as more EVs are deployed, the total EV and CNG miles increase but total CV miles decrease. Figure 6f shows the breakdown of the overall regionwide fleet into REEVs and CVs. Perhaps not surprising, more REEVs are

utilized as they are made available. The BEV case had the same trend with similar magnitudes, and we omit the corresponding figure.

The average regionwide total number of vehicles (regardless of type) across all scenarios was found to be 6,076 and was stable across scenarios. The average number of packages delivered and VHT per vehicle were found to be 99.8 and 9.6 hours, respectively. These statistics show that hours-of-work or duration limits were mostly binding whereas vehicle capacity utilization was roughly 83.2% and hence mostly nonbinding. The average VMT per vehicle was around 55.8 miles with a road travel speed of 35.8 mph, whereas the overall speed (VMT divided by VHT including service times) was 5.8 mph. As a partial validation of our experimental design, we note that similar values have been reported for the average VMT (41.4 miles) of FedEx routes in the United States [Barnitt, 2011, Feng and Figliozzi, 2013]. These statistics are crucial because they reveal potential areas of improvement on which manufacturers and policy makers can focus their attention.

6.4 Impact of Extended Work Hours on Systemwide Performance Metrics

By considering a baseline of 25 EVs per depot, we investigated the systemwide benefits of extended work hours in both the REEV and BEV cases. We focused on five scenarios, with T ranging from 10 (set as a baseline) to 14 hours. As we discussed in the preceding subsection, since vehicle routes are mostly constrained by their work hours (as opposed to their capacity), increasing driving time has an obvious positive impact on performance metrics.

Figure 7 illustrates the regionwide results. First, Figure 7a shows a potential reduction in overall energy costs, of at least 17% and ranging up to 32%, compared with the All CV scenario. Second, Figure 7b and Figure 7c illustrate that both VMT and VHT values decrease strongly as a function of increasing T . Since we also observe a decline in the number of vehicles (see Figure 7f), it is evident that vehicle utilization increases with increasing T . For example, when $T = 14$ hours, the average number of deliveries increases to 115.9, which translates to roughly 96.6% capacity utilization. Figure 7d and Figure 7e break down the VMT by drive mode. We observe that as T is increased, cheaper fuel options (i.e., electricity and CNG) are utilized slightly more while the expensive CV usage decreases. Although not shown, we also observed that vehicle capacity constraints became binding for more routes as work hours were increased.

6.5 Impact of EV Range on Systemwide Performance Metrics

Here we investigate the impact of increasing the EV range in both the REEV and BEV cases. In the REEV case we consider $D_E \in \{33, 40, 60, 80, 100, 125, 150\}$ miles, whereas in the BEV case we consider $D_E \in \{100, 125, 150, 200, 250\}$ miles, where the lower limit of 100 miles is inspired by recent news reports of Abt [2018], Kane [2018].

Figure 8 summarizes the relevant findings. Figure 8a reports cost decreases relative to the All CV scenario. We observe that the cost decrease plateaus in the BEV case after $D_E = 200$ miles. This indicates that very few vehicle routes require traveling more than 200 miles. Moreover, Figure 8b and Figure 8c show that the total VMT (summed across conventional and EV modes) do not change with increasing D_E ; indeed, the average VMT per vehicle remains 55.8 miles across all D_E values, similar to Section 6.3. Therefore, the cost decrease cannot be attributed to a decrease in VMT but rather to a shift from fossil fuel to electricity. This highlights a key takeaway: Although e-commerce LSPs can expect operational cost reductions of roughly 18% by adopting REEVs with an EV range of up to $D_E = 50$ miles, doubling this range further to (very optimistic) values beyond 100 miles will lead to further reductions of at most 5%. Indeed, the average VMT of roughly 56 miles can serve as a good design principle to guide investment in more efficient hybrid powertrains. We note that it is difficult to draw a fair comparison between the REEV and BEV cases, given significant differences in their powertrain technologies and EV ranges, as well as fixed investment and ownership-related costs.

6.6 Impact of Vehicle Capacities on Systemwide Performance Metrics

Compared with the baseline value, we allowed higher vehicle capacities of $Q \in \{150, 180, 210, 240\}$ packages in both the REEV and BEV cases. However, these did not have any significant impacts. Indeed, as noted

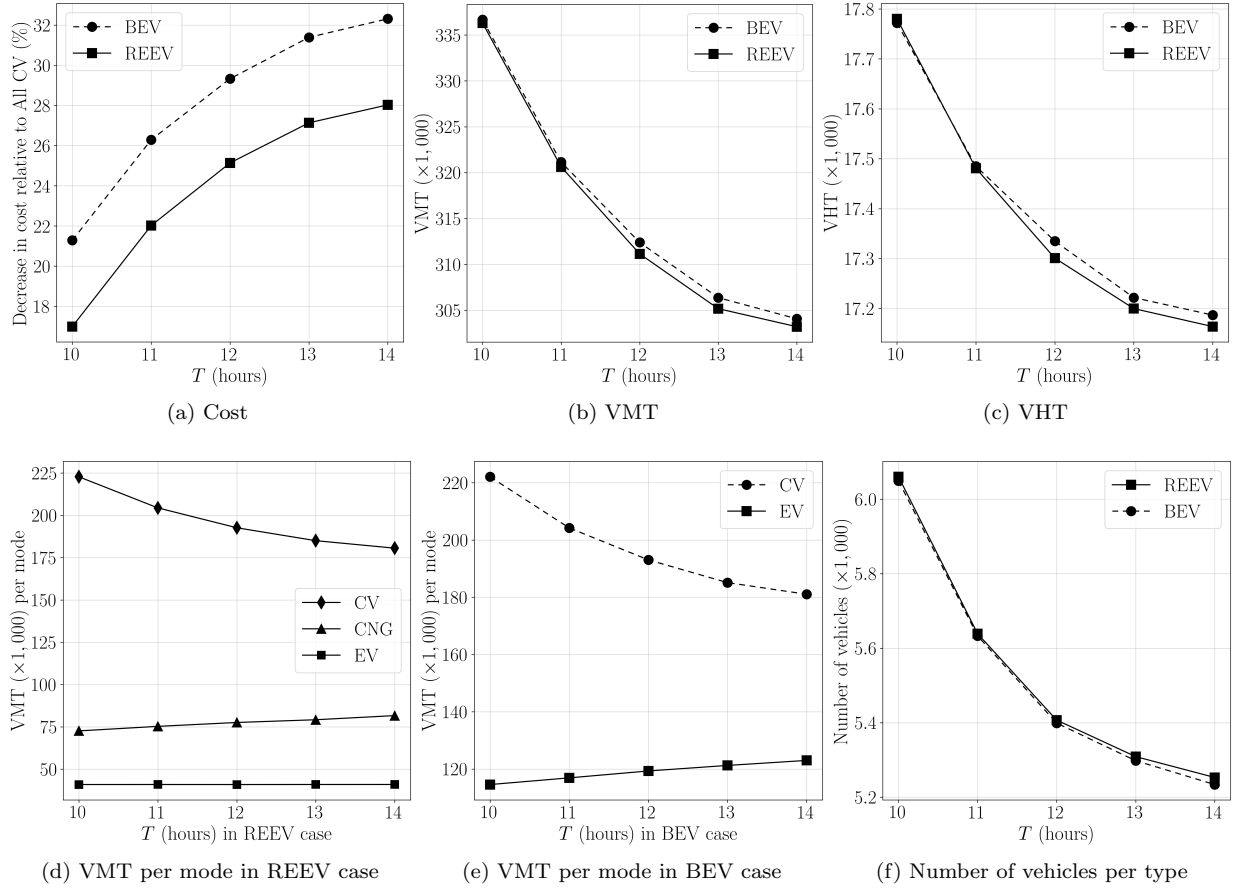


Figure 7: Impact of extended work hours on cost, VMT, VHT, and fleet composition.




Figure 8: Impact of extended EV range on cost and VMT.

in Section 6.3, the average number of packages delivered was roughly 99.8. Therefore, increasing vehicle capacities beyond 120 packages is not expected to yield any benefits. Nonetheless, since the routes are mostly time-limited, it is possible that increasing vehicle capacity while simultaneously reducing the average service time to $\sigma = 2$ minutes (say) could potentially yield cost savings. We investigate this case next.

6.7 Impact of Service Time on Systemwide Performance mMetrics

The choice of the service time is an important parameter that dictates the average route lengths and the number of customers served. We investigate $\sigma \in \{0, 1, 2, 3, 4, 5\}$ minutes focusing only on the REEV case. Note that very low service times are difficult to achieve in practice; we consider them only to establish theoretical upper bounds on savings. Figure 9 summarizes the performance metrics results for these scenarios, relative to a baseline of $\sigma = 5$ minutes. Figure 9a and Figure 9e show that service time alone contributes to roughly 26% of the overall costs and 29.2% of the fleet size in the baseline scenario (compare with $\sigma = 0$). Figure 9b and Figure 9c indicate that VMT and VHT are also significantly impacted by service times. Whereas the total VHT increases roughly linearly with σ , Figure 9d shows that the VHT on major roads (which we define as travel between superlocations) increases almost exponentially with σ . Figure 9f shows that the number of REEVs remains fairly constant across the range of σ values (since all available REEVs are used in each scenario), whereas the number of CVs increases sharply to respond to service time increases.

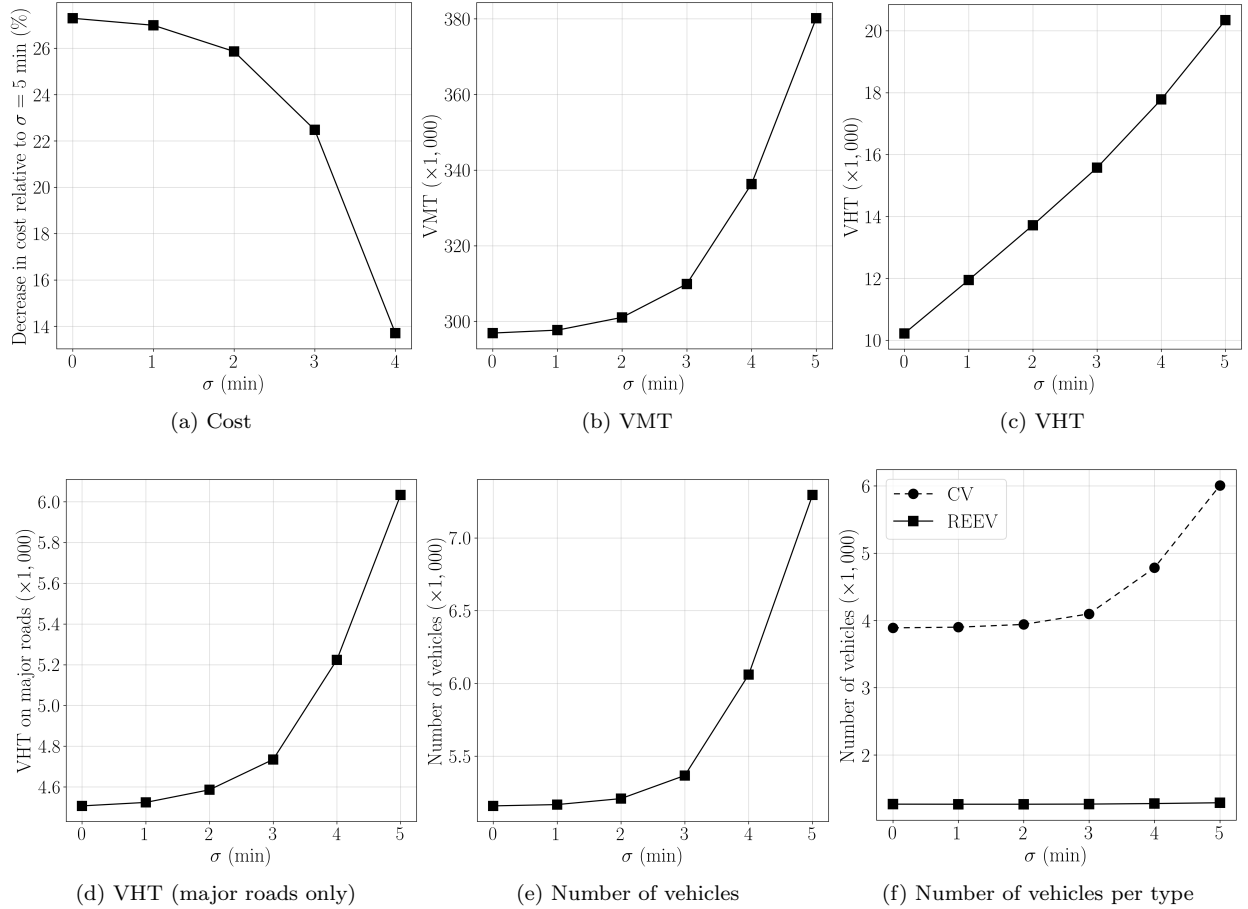


Figure 9: Impact of service time on cost, VMT, VHT, and fleet composition.

7 Conclusions

Range extenders have an important role to play in striking the right balance between the federal government's goal of increasing EV adoption for e-commerce distribution on the one hand and logistics service providers' requirements of achieving operational cost savings on the other. The reason is that by extending the driving range of traditional battery-operated vehicles, they combine the higher efficiency and environmental benefits provided by the latter with the longer range and autonomy of conventional fossil-fuel-powered vehicles. This

combination makes them particularly suited for e-commerce deliveries in dense urban areas, where battery recharging along routes can be too time-consuming to economically justify the use of all-electric vehicles.

This paper investigated the impacts of adopting REEV technology in a large-scale real-world case study concerning parcel deliveries in the Chicago metropolitan area. With a goal to quantify various operational metrics including energy costs, vehicle miles traveled in electric and fossil fuel modes, vehicle hours traveled, and their capacity utilization, we defined a new VRP variant—called the range-extended electric vehicle routing problem, or REEVRP—that simultaneously determines the optimal fleet composition and truck routes for a mixed fleet of REEVs and CVs. Unlike other heterogeneous VRP variants, the REEVRP features heterogeneity in alternative propulsion modes within the same vehicle type. To solve the REEVRP, we proposed an exact BPC algorithm and an ITS metaheuristic. The latter is capable of providing solutions for the large-scale instances in our study featuring more than 3,000 delivery locations, whereas the former provides near-optimal solutions (roughly 5% on average) for small-scale instances. We found that distributors can expect energy cost savings of up to 17% while incurring less than 0.5% increase in VMT and VHT by deploying only a small number of REEVs with a modest EV range of 33 miles. Moreover, since the routes are time-limited, doubling the EV range to 60 miles further reduced costs by only 4%, something that can also be achieved by decreasing the average service time by 1 minute or increasing driver working time by 1 hour. Our case study also revealed other potential areas of improvement to which manufacturers, logistics service providers, and policy makers can focus their attention.

Future directions of study include the design of a more efficient exact BPC algorithm (for example, where the limited EV range of REEVs are directly handled as part of the pricing routine); the modeling of time-dependent travel times and intraroute power management decisions; and sensitivity analyses of parameters that our work did not consider, such as cost coefficients as well as neighborhood-level travel times and distances.

Acknowledgments

This material is based upon work supported by the U.S. Department of Energy, Office of Science, under contract number DE-AC02-06CH11357.

References

- The White House. Fact Sheet: President Biden Announces Steps to Drive American Leadership Forward on Clean Cars and Trucks, 8 2021. Available at <https://www.whitehouse.gov/briefing-room/statements-releases/2021/08/05/fact-sheet-president-biden-announces-steps-to-drive-american-leadership-forward-on-clean-cars-and-trucks/> (accessed DEcember 23, 2021).
- Jacob Bogage. Biden’s spending package would give USPS \$6 billion to replace dangerous mail trucks with electric vehicles, 2021. Available at <https://www.washingtonpost.com/business/2021/10/29/biden-usps-dejoy-trucks/> (accessed DEcember 23, 2021).
- U.S.-China Clean Energy Research Center. Truck Research Utilizing Collaborative Knowledge, 2021. Available at <https://cerc-truck.anl.gov/> (accessed DEcember 23, 2021).
- Kristal Ferchau. Efficient Drivetrains, Inc. delivers industry-first PHEV, renewable CNG armored truck, 2017. Available at <https://www.businesswire.com/news/home/20170406005703/en/Efficient-Drivetrains-Inc.-Delivers-Industry-First-PHEV-Renewable-CNG-Armored-Truck> (accessed DEcember 23, 2021).
- Jacopo Guanetti, Simone Formentin, and Sergio M Savaresi. Energy management system for an electric vehicle with a rental range extender: A least costly approach. *IEEE Transactions on Intelligent Transportation Systems*, 17(11):3022–3034, 2016. doi: 10.1109/TITS.2016.2529962.
- Srikant Jayanthan. Efficient Drivetrains delivers class-6 PHEV with CNG range-extender, 2017. Available at <https://autotechinsight.ihsmarkit.com/news/5239491/efficient-drivetrains-delivers-class-6-phev-with-cng-range-extender> (accessed DEcember 23, 2021).

- Green Car Congress. Hyliion launches Hypertruck Electric Range Extender; CNG genset; 1,000-truck pre-order, 2020. Available at <https://www.greencarcongress.com/2020/06/20200626-hyliion.html> (accessed December 23, 2021).
- Deepak Yadav, Pratchi Singh, and R Prasad. Advanced thermally stable, self-sustaining NiCo_2O_4 catalyst for CNG emissions in lean burn environment. *International Journal of Hydrogen Energy*, 44(54):29057–29065, 2019. doi: 10.1016/j.ijhydene.2019.08.238.
- John Goreham. Why a compressed natural gas plug-in electric hybrid makes sense, 2014. Available at <https://www.torquenews.com/1083/why-compressed-natural-gas-plug-electric-hybrid-makes-sense> (accessed December 23, 2021).
- Jihye Byun and Kyunghwan Choi. Effects analysis of light-duty diesel truck hybrid conversion depending on driving style. *Transportation Research Part D: Transport and Environment*, 97:102958, 2021. doi: 10.1016/j.trd.2021.102958.
- Paolo Toth and Daniele Vigo, editors. *Vehicle Routing: Problems, Methods and Applications*. Society for Industrial and Applied Mathematics, 2014. doi: 10.1137/1.9781611973594.
- Arild Hoff, Henrik Andersson, Marielle Christiansen, Geir Hasle, and Arne Løkketangen. Industrial aspects and literature survey: Fleet composition and routing. *Computers & Operations Research*, 37(12):2041–2061, 2010. doi: 10.1016/j.cor.2010.03.015.
- Stefan Irnich, Michael Schneider, and Daniele Vigo. Four variants of the vehicle routing problem. In Paolo Toth and Daniele Vigo, editors, *Vehicle Routing: Problems, Methods and Applications*, pages 241–271. Society for Industrial and Applied Mathematics, 2014. doi: 10.1137/1.9781611973594.ch9.
- Çağrı Koç, Tolga Bektaş, Ola Jabali, and Gilbert Laporte. Thirty years of heterogeneous vehicle routing. *European Journal of Operational Research*, 249(1):1–21, 2016. doi: 10.1016/j.ejor.2015.07.020.
- Samuel Pelletier, Ola Jabali, and Gilbert Laporte. 50th anniversary invited article – goods distribution with electric vehicles: Review and research perspectives. *Transportation Science*, 50(1):3–22, 2016. doi: 10.1287/trsc.2015.0646.
- Tomislav Erdelić and Tonči Carić. A survey on the electric vehicle routing problem: Variants and solution approaches. *Journal of Advanced Transportation*, 2019, 2019. doi: <https://doi.org/10.1155/2019/5075671>.
- Ryan G Conrad and Miguel Andres Figliozzi. The recharging vehicle routing problem. In *Proceedings of the 2011 Industrial Engineering Research Conference*, pages 2785–2792, 2011.
- Sevgi Erdoğan and Elise Miller-Hooks. A green vehicle routing problem. *Transportation Research Part E: Logistics and Transportation Review*, 48(1):100–114, 2012. doi: 10.1016/j.tre.2011.08.001.
- Michael Schneider, Andreas Stenger, and Dominik Goeke. The electric vehicle-routing problem with time windows and recharging stations. *Transportation Science*, 48(4):500–520, 2014. doi: 10.1287/trsc.2013.0490.
- Ángel Felipe, M Teresa Ortuño, Giovanni Righini, and Gregorio Tirado. A heuristic approach for the green vehicle routing problem with multiple technologies and partial recharges. *Transportation Research Part E: Logistics and Transportation Review*, 71:111–128, 2014. doi: 10.1016/j.tre.2014.09.003.
- Alejandro Montoya, Christelle Guéret, Jorge E Mendoza, and Juan G Villegas. A multi-space sampling heuristic for the green vehicle routing problem. *Transportation Research Part C: Emerging Technologies*, 70:113–128, 2016. doi: 10.1016/j.trc.2015.09.009.
- Guy Desaulniers, Fausto Errico, Stefan Irnich, and Michael Schneider. Exact algorithms for electric vehicle-routing problems with time windows. *Operations Research*, 64(6):1388–1405, 2016. doi: <https://doi.org/10.1287/opre.2016.1535>.

- Merve Keskin and Bülent Çatay. Partial recharge strategies for the electric vehicle routing problem with time windows. *Transportation Research Part C: Emerging Technologies*, 65:111–127, 2016. doi: 10.1016/j.trc.2016.01.013.
- Aurélien Froger, Jorge E Mendoza, Ola Jabali, and Gilbert Laporte. Improved formulations and algorithmic components for the electric vehicle routing problem with nonlinear charging functions. *Computers & Operations Research*, 104:256–294, 2019. doi: 10.1016/j.cor.2018.12.013.
- Angel A Juan, Jarrod Goentzel, and Tolga Bektas. Routing fleets with multiple driving ranges: Is it possible to use greener fleet configurations? *Applied Soft Computing*, 21:84–94, 2014. doi: 10.1016/j.asoc.2014.03.012.
- Ons Sassi, Wahiba Ramdane Cherif-Khettaf, and Ammar Oulamara. Vehicle routing problem with mixed fleet of conventional and heterogeneous electric vehicles and time dependent charging costs. *International Journal of Mathematical and Computational Sciences*, 9(3):171–181, 2015. doi: 10.5281/zenodo.1099756.
- Dominik Goeke and Michael Schneider. Routing a mixed fleet of electric and conventional vehicles. *European Journal of Operational Research*, 245(1):81–99, 2015. doi: 10.1016/j.ejor.2015.01.049.
- Philippe Lebeau, Cedric De Cauwer, Joeri Van Mierlo, Cathy Macharis, Wouter Verbeke, and Thierry Coosemans. Conventional, hybrid, or electric vehicles: Which technology for an urban distribution centre? *The Scientific World Journal*, 2015, 2015. doi: 10.1155/2015/302867.
- Gerhard Hiermann, Jakob Puchinger, Stefan Ropke, and Richard F Hartl. The electric fleet size and mix vehicle routing problem with time windows and recharging stations. *European Journal of Operational Research*, 252(3):995–1018, 2016. doi: 10.1016/j.ejor.2016.01.038.
- Samuel Pelletier, Ola Jabali, and Gilbert Laporte. Charge scheduling for electric freight vehicles. *Transportation Research Part B: Methodological*, 115:246–269, 2018. doi: 10.1016/j.trb.2018.07.010.
- Alexandre M Florio, Nabil Absi, and Dominique Feillet. Routing electric vehicles on congested street networks. *Transportation Science*, 55(1):238–256, 2021. doi: 10.1287/trsc.2020.1004.
- Tarek Abdallah. The Plug-In Hybrid Electric Vehicle Routing Problem with Time Windows. Master’s thesis, University of Waterloo, 2013. <http://hdl.handle.net/10012/7582>.
- Simona Mancini. The hybrid vehicle routing problem. *Transportation Research Part C: Emerging Technologies*, 78:1–12, 2017. doi: 10.1016/j.trc.2017.02.004.
- Gerhard Hiermann, Richard F Hartl, Jakob Puchinger, and Thibaut Vidal. Routing a mix of conventional, plug-in hybrid, and electric vehicles. *European Journal of Operational Research*, 272(1):235–248, 2019. doi: 10.1016/j.ejor.2018.06.025.
- F Yu Vincent, AAN Perwira Redi, Yosi Agustina Hidayat, and Oktaviyanto Jimat Wibowo. A simulated annealing heuristic for the hybrid vehicle routing problem. *Applied Soft Computing*, 53:119–132, 2017. doi: 10.1016/j.asoc.2016.12.027.
- Xiaohui Li, Xuemin Shi, Yi Zhao, Huagang Liang, and Yuan Dong. SVND enhanced metaheuristic for plug-in hybrid electric vehicle routing problem. *Applied Sciences*, 10(2):441, 2020. doi: 10.3390/app10020441.
- Christian Doppstadt, Achim Koberstein, and Daniele Vigo. The hybrid electric vehicle–traveling salesman problem. *European Journal of Operational Research*, 253(3):825–842, 2016. doi: <https://doi.org/10.1016/j.ejor.2016.03.006>.
- Sina Bahrami, Mehdi Nourinejad, Glareh Amirjamshidi, and Matthew J Roorda. The plugin hybrid electric vehicle routing problem: A power-management strategy model. *Transportation Research Part C: Emerging Technologies*, 111:318–333, 2020. doi: 10.1016/j.trc.2019.12.006.

- Giovanni De Nunzio, Ibtiel Ben Gharbia, and Antonio Sciarretta. A general constrained optimization framework for the eco-routing problem: Comparison and analysis of solution strategies for hybrid electric vehicles. *Transportation Research Part C: Emerging Technologies*, 123:102935, 2021. doi: 10.1016/j.trc.2020.102935.
- Adrian Casparia, Steffen Fahra, and Alexander Mitsos. Optimal Eco-Routing for Hybrid Vehicles with Mechanistic/Data-Driven Powertrain Model Embedded, 2021. Available on Optimization Online. http://www.optimization-online.org/DB_HTML/2021/06/8444.html.
- Joshua Auld, Michael Hope, Hubert Ley, Vadim Sokolov, Bo Xu, and Kuilin Zhang. POLARIS: Agent-based modeling framework development and implementation for integrated travel demand and network and operations simulations. *Transportation Research Part C: Emerging Technologies*, 64:101–116, 2016. doi: 10.1016/j.trc.2015.07.017.
- Chi-Kin Chau, Khaled Elbassioni, and Chien-Ming Tseng. Drive mode optimization and path planning for plug-in hybrid electric vehicles. *IEEE Transactions on Intelligent Transportation Systems*, 18(12):3421–3432, 2017. doi: 10.1109/tits.2017.2691606.
- Marcus Poggi and Eduardo Uchoa. New exact algorithms for the capacitated vehicle routing problem. In Paolo Toth and Daniele Vigo, editors, *Vehicle Routing: Problems, Methods, and Applications*, pages 59–86. Society for Industrial and Applied Mathematics, 2014. doi: 10.1137/1.9781611973594.ch3.
- Roberto Baldacci, Aristide Mingozzi, and Roberto Roberti. New route relaxation and pricing strategies for the vehicle routing problem. *Operations Research*, 59(5):1269–1283, 2011. doi: 10.1287/opre.1110.0975.
- Diego Pecin, Artur Pessoa, Marcus Poggi, and Eduardo Uchoa. Improved branch-cut-and-price for capacitated vehicle routing. *Mathematical Programming Computation*, 9(1):61–100, 2017a. doi: 10.1007/s12532-016-0108-8.
- Artur Pessoa, Ruslan Sadykov, Eduardo Uchoa, and François Vanderbeck. A generic exact solver for vehicle routing and related problems. *Mathematical Programming*, 183(1):483–523, 2020. doi: 10.1007/s10107-020-01523-z.
- Luigi Di Puglia Pugliese and Francesca Guerriero. A survey of resource constrained shortest path problems: Exact solution approaches. *Networks*, 62(3):183–200, 2013. doi: 10.1002/net.21511.
- Ruslan Sadykov, Eduardo Uchoa, and Artur Pessoa. A bucket graph-based labeling algorithm with application to vehicle routing. *Transportation Science*, 55(1):4–28, 2021. doi: 10.1287/trsc.2020.0985.
- Ricardo Fukasawa, Humberto Longo, Jens Lysgaard, Marcus Poggi De Aragão, Marcelo Reis, Eduardo Uchoa, and Renato F Werneck. Robust branch-and-cut-and-price for the capacitated vehicle routing problem. *Mathematical Programming*, 106(3):491–511, 2006. doi: 10.1007/s10107-005-0644-x.
- Giovanni Righini and Matteo Salani. New dynamic programming algorithms for the resource constrained elementary shortest path problem. *Networks: An International Journal*, 51(3):155–170, 2008. doi: 10.1002/net.20212.
- Stefan Irnich, Guy Desaulniers, Jacques Desrosiers, and Ahmed Hadjar. Path-reduced costs for eliminating arcs in routing and scheduling. *INFORMS Journal on Computing*, 22(2):297–313, 2010. doi: 10.1287/ijoc.1090.0341.
- Artur Pessoa, Ruslan Sadykov, Eduardo Uchoa, and François Vanderbeck. Automation and combination of linear-programming based stabilization techniques in column generation. *INFORMS Journal on Computing*, 30(2):339–360, 2018. doi: 10.1287/ijoc.2017.0784.
- Roberto Roberti and Aristide Mingozzi. Dynamic ng-path relaxation for the delivery man problem. *Transportation Science*, 48(3):413–424, 2014. doi: 10.1287/trsc.2013.0474.

- Mads Jepsen, Bjørn Petersen, Simon Spoorendonk, and David Pisinger. Subset-row inequalities applied to the vehicle-routing problem with time windows. *Operations Research*, 56(2):497–511, 2008. doi: 10.1287/opre.1070.0449.
- Diego Pecin, Claudio Contardo, Guy Desaulniers, and Eduardo Uchoa. New enhancements for the exact solution of the vehicle routing problem with time windows. *INFORMS Journal on Computing*, 29(3):489–502, 2017b. doi: 10.1287/ijoc.2016.0744.
- Roberto Baldacci, Nicos Christofides, and Aristide Mingozzi. An exact algorithm for the vehicle routing problem based on the set partitioning formulation with additional cuts. *Mathematical Programming*, 115(2):351–385, 2008. doi: 10.1007/s10107-007-0178-5.
- Puca Huachi Vaz Penna, Anand Subramanian, and Luiz Satoru Ochi. An iterated local search heuristic for the heterogeneous fleet vehicle routing problem. *Journal of Heuristics*, 19(2):201–232, 2013. doi: 10.1007/s10732-011-9186-y.
- Anirudh Subramanyam, Panagiotis P Repoussis, and Chrysanthos E Gounaris. Robust optimization of a broad class of heterogeneous vehicle routing problems under demand uncertainty. *INFORMS Journal on Computing*, 32(3):661–681, 2020. doi: 10.1287/ijoc.2019.0923.
- Helena R Lourenço, Olivier C Martin, and Thomas Stützle. Iterated local search: Framework and applications. In Michel Gendreau and Jean-Yves Potvin, editors, *Handbook of Metaheuristics*, pages 363–397. Springer US, 2010. doi: 10.1007/978-1-4419-1665-5_12.
- Emile Aarts and Jan Karel Lenstra. *Local Search in Combinatorial Optimization*. Princeton University Press, 2003. doi: 10.2307/j.ctv346t9c.
- Fred Glover. Tabu search and adaptive memory programming – Advances, applications and challenges. In Richard S Barr, Richard V Helgason, and Jeffery L Kennington, editors, *Interfaces in Computer Science and Operations Research: Advances in Metaheuristics, Optimization, and Stochastic Modeling Technologies*, pages 1–75. Springer, 1997. doi: 10.1007/978-1-4615-4102-8_1.
- Birger Funke, Tore Grünert, and Stefan Irnich. Local search for vehicle routing and scheduling problems: Review and conceptual integration. *Journal of Heuristics*, 11(4):267–306, 2005. doi: 10.1007/s10732-005-1997-2.
- Federal Highway Administration. National household travel survey, 2017. Available at <https://nhts.ornl.gov> (accessed December 23, 2021).
- John Spadafora and Marifer Rodriguez. Pitney Bowes parcel shipping index reveals 37 percent parcel volume growth in US for 2020, 2021. Available at <https://www.businesswire.com/news/home/20210914005274/en/Pitney-Bowes-Parcel-Shipping-Index-Reveals-37-Percent-Parcel-Volume-Growth-in-US-for-2020> (accessed December 23, 2021).
- U.S. Energy Information Administration. Table 5.3. Average Price of Electricity to Ultimate Customers, 2021. Available at https://www.eia.gov/electricity/monthly/epm_table_grapher.php?t=epmt_5_03 (accessed December 23, 2021).
- Alternative Fuels Data Center. National Average Price Between July 1 and July 15, 2021, 2021. Available at <https://afdc.energy.gov/fuels/prices.html> (accessed December 23, 2021).
- M Lammert. Twelve-month evaluation of UPS diesel hybrid electric delivery vans. Technical report, National Renewable Energy Laboratory, Golden, CO, 2009. URL <https://www.nrel.gov/docs/fy10osti/44134.pdf>. No. NREL/TP-540-44134.
- Robb Barnitt. FedEx express gasoline hybrid electric delivery truck evaluation: 12-month report. Technical report, National Renewable Energy Lab, Golden, CO, United States, 2011.
- IBM. IBM ILOG CPLEX Optimizer, 2021. Available at <https://www.ibm.com/analytics/cplex-optimizer> (accessed December 23, 2021).

Wei Feng and Miguel Figliozzi. An economic and technological analysis of the key factors affecting the competitiveness of electric commercial vehicles: A case study from the USA market. *Transportation Research Part C: Emerging Technologies*, 26:135–145, 2013. doi: 10.1016/j.trc.2012.06.007.

Neil Abt. UPS working with Thor on electric class 6 truck, 2018. Available at <https://www.fleetowner.com/equipment/powertrain/article/21702782/ups-working-with-thor-on-electric-class-6-truck> (accessed December 23, 2021).

Mark Kane. UPS teams with Thor for electric trucks, 2018. Available at <https://insideevs.com/news/338598/ups-teams-with-thor-for-electric-trucks/> (accessed December 23, 2021).

The submitted manuscript has been created by UChicago Argonne, LLC, Operator of Argonne National Laboratory (“Argonne”). Argonne, a U.S. Department of Energy Office of Science laboratory, is operated under Contract No. DE-AC02-06CH11357. The U.S. Government retains for itself, and others acting on its behalf, a paid-up nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan. <http://energy.gov/downloads/doe-public-access-plan>.