

Routing and Scheduling an Electric Vehicle Fleet Handling Dynamic Customer Requests

Panagiotis Liakos, Iosif Angelidis, and Alex Delis

University of Athens, GR15703, Athens, Greece
{p.liakos,sdi1100127,ad}@di.uoa.gr

Abstract. Ongoing worldwide environmental issues along with often-voiced consumer concerns have paved the way so that governments legislate and help usher into operation alternative-fueled vehicles and pertinent infrastructures. In this regard, battery-powered electric vehicles have been introduced in the last decade and the service industry has been deploying such trucks in their distribution networks. However, electric vehicles do impose limitations as far as their traveling range is concerned. Moreover, replenishing the power to the vehicle batteries may entail lengthy charging visits at respective stations. In this paper, we examine the problem of routing and scheduling a fleet of electric vehicles that seek to satisfy dynamic pickup and delivery requests in an urban environment. We develop a web application to facilitate cooperation between organizations and individuals involved in urban freight transport. Geolocation services and mobile devices enable us to manage the fleet and make timely decisions. Moreover, we propose three heuristic-based recharging strategies to ensure that electric vehicles can restore their energy levels in an effective manner. Through detailed simulation experiments, we show that the costs associated with the use of an electric vehicle fleet concern mainly the size of the fleet. The impact regarding the total route length traveled is less evident for all our strategies.

Keywords: Urban Computing; Electric Vehicles; GNSS

1 Introduction

The always-timely concern for the impact that transportation and logistics have on the environment, has played a pivotal role in shifting the attention of governments, commercial actors as well as consumers to practices collectively known as *green logistics*. The aim of green logistics is to not only consider economic factors, but also environmental and social aspects so that communities can progressively attain production and distribution of goods in a sustainable way [28]. It has been also established that the cost for achieving noteworthy reductions in greenhouse effects is by and large modest, especially if such costs are amortized over time [7]. As central and regional governments enforce adopted environmental regulations and citizens call for the embrace of innovative means for transporting goods, companies and organizations have been responding by becoming early adopters

of novel green practices [17]. Among all those involved, companies whose main line of business is in transportation and logistics have respectively shown great interest and have considered the wide-spread use of *Electric Vehicles (EVs)* to comply with environmental state legislation and pertinent city ordinances. As a matter of fact, a number of companies in the small-package shipping industry including *DHL*, *DPD* and *UPS*, have been reportedly using electric vehicles for last-mile delivery for some time already [16].

Logistics companies do face numerous challenges when it comes to fulfilling their distribution tasks and attempting to meet deadlines. Studies indicate that freight vehicles represent no more than 15% of total traffic flow in urban areas [1], but due to their size and frequent stops for deliveries have a more significant green-house impact than passenger vehicles. Moreover, diesel-powered freight vehicles generate emissions that are very harmful to people. The total carbon dioxide emitted by all forms of transport in London in 2006 was 9.6 tonnes, of which an impressive 23% was produced by freight vehicles alone [2]. The significant threat to public health and safety as well as the negative impact that such emissions have on climate change could be addressed by using vehicles running on alternative fuel. A further step would be to ideally cut down on the number of vehicles traveling about urban areas by possibly increasing the refueling rate of trucks already set in motion. Moreover, traffic congestion, air pollution and unnecessary costs could all be reduced by minimizing the total mileage needed to carry out by the fleet of the vehicles involved in the distribution roster.

The numerous and diverse aspects and tasks entailed in the transportation and distribution of goods are typically modeled using one of the many generalizations of the classic *NP-hard vehicle routing problem (VRP)* [6]. The main concern here is to achieve minimization of the total costs incurred during the transportation. A well studied flavor of the VRP is the *Dynamic Pickup and Delivery Problem with Time Windows (PDPTW)* [24]. Advances in *Intelligent Transportation Systems (ITSs)* including geolocation and object-tracking have enabled the use of techniques to address the PDPTW problem in real-world operational settings. This has been the case with companies offering same-day pickup and delivery of letters and parcels. As devices and tools become amply available with the introduction of the *Internet of Things*, the applicability of the PDPTW-problem is on the rise. Among the applications that seize such opportunities for commercial success, *UberPool* [32] is likely the most prominent example as it manages a fleet of vehicles serving simultaneously multiple requests for pickups and deliveries of passengers.

Previous approaches on the Dynamic PDPTW ([24,25,9,26,21,20,3]) predominantly focus on conventional vehicles and do not address the issue of scheduling refueling stops, as the corresponding time needed is negligible. However, electric vehicles (EVs) offer limited range which, at the moment, is not sufficient for the typical delivery tours of logistics and transportation service providers [29]. Consequently, EVs used in such context will require to visit recharging stations along their route(s) to replenish their power supply. We should point out that the recharging times needed by EVs might be fairly significant, if compared to that

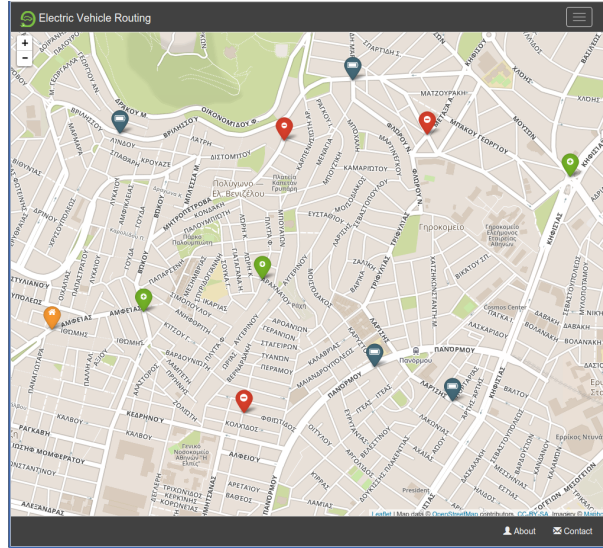


Fig. 1. Pickup (green) and delivery (red) requests serviced through a fleet of electric vehicles (orange). Recharging stations (blue) allow vehicles to restore their energy levels but cost in terms of time and route length. The objective is to assign the requests to vehicles in an effort to minimize the size of the fleet and the route length.

of conventional vehicles. The use of a fleet of EVs does evidently add another dimension to the PDPTW problem, by incorporating the possibility of stopping for a recharge using an appropriate policy. This is evident in Figure 1, which illustrates three client requests (pickups and deliveries) and four recharging stations. If conventional vehicles were used, the problem of routing and scheduling these requests would be to involve the minimum amount of vehicles and assign the requests to those that minimize the total route length. With the appearance of EVs and their inherent refueling issues, the logistics problem at hand is certainly more complicated as we now have to cater for lengthier refueling stops. Gonçalves et al. [14] approached the issue at hand by considering a mixed fleet of conventional and electric vehicles. However, their model does not consider specific locations for the recharging stations. Instead, recharges are possible anywhere, which might not be the case for quite a while into the future.

In this paper, we develop a framework for the effective coordination and monitoring of the parties involved in urban distribution using electric vehicles. We consider a fleet of EVs equipped with GNSS receivers that satisfies real-time customer requests submitted through a REST API. Recharging stations are positioned in specific locations and EVs may visit them while executing their schedules to perform needed power refueling. Recharging times depend on the charging level of vehicles. Furthermore, we formally introduce the *Electric Vehicle Dynamic Pickup and Delivery Problem with Time Windows and Recharging Stations (EV-DPDPTW)* and develop an on-line algorithm to find

approximate solutions. We adopt the *rolling horizon* principle and the *drive first* waiting strategy respectively described in [24] and [21], to address the scheduling and routing of client requests. In addition, we propose three heuristic-based strategies that help address the problem of recharging as currently experienced by EVs and examine both their pros and cons. The **Eager Recharging** strategy exploits every opportunity of performing a recharge by visiting all stations that are *close* to client requests, regardless of the vehicle’s energy level. The **Lazy Recharging** strategy awaits until the battery of each vehicle is close to exhaustion before scheduling a visit to a recharging station. Finally, our proposed **Smart Recharging** technique offers a hybrid of the two prior approaches as it schedules visits to nearby recharging stations before battery exhaustion, provided that the energy level is below a preset level. We show that the new dimension of the examined EV-DPDPWTW problem, i.e., the need for lengthy returns to refueling spots, plays a significant role and affects the fleet size regardless of the strategy followed. Through experimentation, we also show that the **Smart Recharging** strategy does indeed benefit from the advantages of the **Eager** and **Lazy Recharging** strategies, and outperforms them with regards to both the fleet size and the route length. Finally, we present results regarding the average energy level at which each strategy schedules visits to recharging stations, a choice that has an impact on the lifespan of a vehicle’s battery.

The rest of this paper is organized as follows: Section 2 provides definitions for the PDPTW problems and requisite notions that are helpful in describing our approach; the section also offers some background information helpful to better understand the issues arising from the introduction of the EV fleet. Section 3 outlines the architecture of our approach and details our proposed three strategies. We then present our simulation-derived results in Section 4. Finally, we review related work and offer concluding remarks and directions for future work in Sections 5 and 6 respectively.

2 Preliminaries

In this section, we review some basic definitions that are helpful in introducing our approach and briefly outline technology aspects related to the deployment of EV fleets. In particular, we place emphasis on vehicles based on batteries.

2.1 The Pickup and Delivery Problem with Time-Windows

We commence by formally outlining the *Static* and *Dynamic* versions of the PDPTW problem and by furnishing key aspects as defined in [20,21].

The Static PDPTW Let $P^+ = \{1^+, 2^+, \dots, n^+\}$ be a set of pickup locations, and $P^- = \{1^-, 2^-, \dots, n^-\}$ a set of corresponding delivery locations. Pairs (i^+, i^-) , where $i^+ \in P^+$ and $i^- \in P^-$, represent transportation requests for performing a pickup at location i^+ and the associated delivery at location i^- . We denote by d_{ij} the distance from location i to location j , by t_{ij} the travel time

from i to j , by s_i the service time at location i , and by $[a_i, b_i]$ the time window of a pickup or delivery request i .

The *Pickup and Delivery Problem with Time-Windows* is about determining a set of optimal routes and corresponding schedules for a fleet of vehicles in order to serve these transportation requests with respect to the following constraints:

1. Each route starts at the corresponding vehicle's embarking position.
2. A pickup and its associated delivery are satisfied by the same vehicle.
3. A pickup is always made before its associated delivery.
4. All time windows are satisfied.
5. A vehicle is allowed to wait at its embarking position or at any pickup or delivery location.
6. The total distance traveled by vehicles is minimized.

A solution to the PDPTW determines an ordered sequence of locations for each vehicle route (routing) and the arrival and departure times for all locations of each route (scheduling). The PDPTW reduces to the Multiple Traveling Salesman Problem with Time Windows when the pickup and delivery locations of every request coincide. The PDPTW is *NP-hard*, and deciding whether there exists a feasible solution when the number of vehicles is fixed is *NP-complete* in the strong sense [27].

The Dynamic PDPTW The *Dynamic* version of the PDPTW drops the assumption that all information is available at the time of planning. Instead, in the Dynamic PDPTW the problem is to allocate the requests in question to vehicles in an online fashion. This setting is much more applicable in real-life situations, where requests are expected to be made dynamically and a priori planning of operations is not feasible. The real-time allocation of customers to vehicles yields an array of additional issues when it comes to scheduling; choices selected by online algorithmic techniques may lead to a total distance that is larger than the one traveled should all requisite information were known in advance. Fortunately, heuristic-based approaches are known to offer good approximate solutions to the Dynamic PDPTW problem [21,20].

2.2 EVs and Battery Developments

EVs are divided into three main types: Battery Electric Vehicles (BEVs), (Plug-in) Hybrid Electric Vehicles (PHEVs, HEVs), and Fuel Cell Electric Vehicles (FCEVs). The cost of FCEVs is considered prohibitive at this moment [22]. In this paper, we concentrate on BEVs which display two favorable features: they do not produce any emissions themselves and cause limited noise while in operation; both of these features are due to the fact that such vehicles base their motion and overall function entirely on batteries. These advantages have led to strong government support on the development of BEV technology, which in turn is advancing at an unprecedented pace.

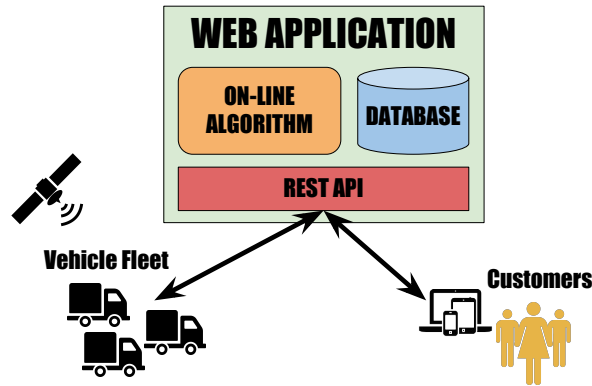


Fig. 2. Architecture of our approach.

The development of commercially viable battery technology has played a major role in the widespread use of EVs. Although battery energy densities are expected to triple by 2030 [8], contemporary batteries are far more inferior if compared to gasoline [15]. This results in noticeably long replenishing times for car batteries and a relatively low driving range [23]. The preferred option for BEVs is lithium-ion batteries, as their energy and power density as well as their battery lifespan are higher in respect to other alternatives, including nickel-metal hydride and the inexpensive lead-acid batteries [18].

When it comes to operation and management of EV commercial fleets, the time needed for re-charging EV batteries is crucial. More specifically, this time period highly depends on the size of the batteries [33] and the type of facility used; refueling of batteries may last from 30 minutes to multiple hours. Achieving a full capacity battery recharge may call for a very long –if not outright excessive– period of waiting time at the re-charging station. The latter is due to the fact that the final phase of a recharge is not linear with respect to time and can itself last for several hours [4]. Vehicle range and overall battery life are influenced by the pattern(s) BEVs are charged and discharged. In particular, frequent discharges to deep levels shorten their lifespan [19]. Similarly, frequently charging close to the maximum capacity rapidly leads to battery deterioration [31]. Such usage patterns do limit the usable battery capacity. BEVs reportedly reach ranges of 150 miles on a single charge, which essentially restricts their usage to urban areas [11]. The typical delivery tours for logistics and transportation service providers do surpass this value [29]. Therefore, visits to recharging stations along routes are deemed necessary in the course of a business day.

3 Overview of our Approach

We present here our framework for managing a fleet of EVs handling dynamic pickup and delivery requests. Our architecture, depicted in Figure 2, features a

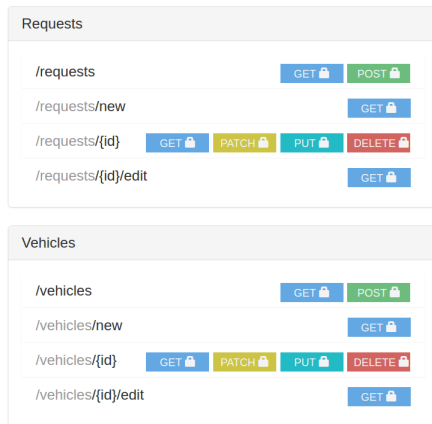


Fig. 3. REST API

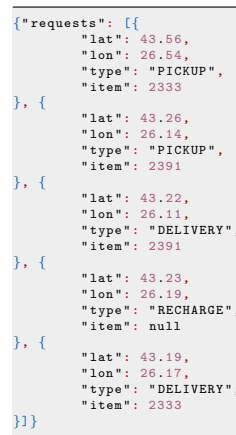


Fig. 4. JSON response

Web Application that employs an on-line algorithm specifically designed for the EV-DPDPTW. A REST API handles communication with vehicles equipped with GNSS receivers and customers that have access through desktops, laptops or mobile devices. Information submitted to the Web Application is persisted to a database.

We present in Figure 3 part of our REST API.¹ Actions listed under ‘Requests’ in Figure 3, enable authenticated customers to submit, edit and view (their) requests. Submitted requests are passed to the on-line algorithm in order to be assigned to a vehicle. Additionally, users with administrative privileges are authorized to monitor the delivery process through the ‘Vehicles’ actions. In particular, issuing a GET ‘/vehicles’ or ‘/vehicles/{id}’ request provides a view of the status of all vehicles or the one specified with the ‘id’ parameter, respectively. We provide illustrations of such views in Figure 5. We can see that OpenStreetMap² powered maps report the position of the vehicle at any time, as well as the customer requests associated with the vehicle. Vehicles are also enabled to communicate with the Web Application through the REST API. Their scheduling information can be retrieved through an authenticated GET ‘/vehicles/{id}’ requesting JSON content. Moreover, vehicles may update their current state through a PATCH ‘/vehicles/{id}’ specifying their location or list of items. An exemplary response to the first request is illustrated in Figure 4. We can see that the vehicle is scheduled to pickup and deliver 2 items requests and will visit a recharging station in between.

Scheduling information is derived from our on-line algorithm for the EV-DPDPTW. We employ the *rolling horizon* principle and the *Drive First* waiting strategy, described in [24] and [21], respectively. Moreover, we propose heuristics to address the additional issues that arise due to the use of BEVs, as well as the

¹ Our REST API is documented using [raml2html](#).

² <http://www.openstreetmap.org/>

Algorithm 1: On-line algorithm for the EV-DPDPTW.

```
input  : Incoming requests  $P^+ \cup P^-$ .  
output : Routing and scheduling information.  
1 begin  
2   while true do  
3     select the eligible unassigned requests;  
4     foreach eligible request  $(i^+, i^-)$  do  
5       foreach available vehicle do  
6         find the best insertion;  
7         if insertion triggers recharge then  
8           schedule recharge visit for the vehicle;  
9       if there are vehicles that can serve  $(i^+, i^-)$  then  
10        select the best insertion;  
11        insert the request in the selected vehicle;  
12      else  
13        assign the request to a new vehicle;  
14      update scheduling information;
```

presence of limited recharging stations. Algorithm 1 provides pseudo-code for our on-line approach, which extends the *cheapest insertion procedure* described in [21] in order to handle visits to recharging stations as well. Our algorithm awaits for incoming requests (2). Once a request arrives, we examine all possible pairs of feasible slots in each route to schedule the request’s pickup and delivery (6). Insertions in these slots may trigger a need for a visit to a recharging station to restore the charging level of the vehicle’s battery (7-8). If there are vehicles that can accommodate this request, the algorithm selects the one that minimizes the total route length (10) and inserts the request (11). If, however, there are no vehicles that can handle the request, we extend our fleet with a new vehicle, and insert the request to it (13). Finally, we update the scheduling information associated with our fleet (14). Lines (7-8) outline the operations needed to ascertain if a recharge is needed and ensure that the vehicles of the fleet remain in operational state with regards to their battery charging level. We note that we consider the uncapacitated version of the problem, that is applicable in cases of carriers transporting small parcels or letters. Our algorithm is easily extensible to the capacitated version simply by adding a corresponding constraint.

In what follows, we define three heuristic-based recharging strategies, namely **Eager**, **Lazy**, and **Smart Recharging**, that we use to decide when to schedule a visit to a recharging station.

3.1 Eager Recharging

The **Eager Recharging** strategy requires a vehicle to perform a recharge after every delivery request it satisfies, given that the delivery location is close to a

recharging station. In particular, having found the best insertion of a request in a route, we additionally investigate the possibility of scheduling a stop to a recharging station by estimating the distance between the delivery location and its closest recharging station. In the case that this distance is smaller than a predefined limit, we additionally schedule the recharging stop. A visit to a recharging station is also scheduled in case the energy level of the vehicle is not sufficient for further client requests.

Figure 5a illustrates a route formed after following the **Eager Recharging** strategy. We observe that our online algorithm schedules visits to recharging stops after two of the three deliveries that are assigned to a particular vehicle, as the distance of the corresponding station from delivery locations is small. The third delivery location does not have a nearby recharging station and the energy level of the vehicle is sufficient for new requests due to the previous recharges.

The intuition behind the **Eager Recharging** strategy is that long recharging stops may have a negative impact on the number of vehicles required to satisfy the customer requests. Recharging of a battery is (up to a point) linear with respect to time [4], and thus, vehicles that fully wind up their batteries need to spend more time in recharging stations. Given the current range of electric vehicles, we expect them to require visits to recharging stations along their routes. Therefore, allowing vehicles to exhaust their energy levels is likely to lead to multiple vehicles being detained simultaneously, and, inevitably, to an unnecessary expansion of the fleet.

Furthermore, the **Eager Recharging** strategy favors the dispersion of vehicles in different locations of an area in a way similar to the *waiting strategy* of Mitrović-Minić and Laporte [21] by extending their stay close to recharging stations. As the case is with the *waiting strategy*, we expect that this will have a negative impact on the size of the fleet, but may result to better vehicle assignments for specific clients with pickups close to a recharging station.

The main drawback of the **Eager Recharging** strategy is that it does not allow a lot of reordering to occur. As new clients arrive, our algorithm examines the possibility of serving them in between requests that have already been scheduled, as long as the time windows are not violated. The **Eager Recharging** strategy leads to multiple small visits to recharging stops which essentially limit the opportunities of re-evaluating scheduling decisions.

3.2 Lazy Recharging

The **Lazy Recharging** strategy requires a vehicle to perform a recharge only in cases when the battery charging level is not enough to service any more incoming requests. As there is a risk of exhausting the battery of a vehicle before ever reaching to a station, this strategy acts proactively. In particular, with **Lazy Recharging** a visit to a recharging station is scheduled at the end of every current route but is consolidated only when there is no eligible request that can feasibly be serviced before the scheduled recharging stop. Until then, the visit is re-scheduled after the latest delivery planned in the route every time a new request is placed in the end of the route.

Figure 5b depicts the impact of the **Lazy Recharging** strategy. We see that although the vehicle is scheduled to pass nearby recharging station locations, it is not scheduled to visit one of them. A visit to a recharging station occurs only when the battery of the vehicle is close to being exhausted, even though there is not a recharging station close by. Moreover, we observe that the order in which the pickups and deliveries occur is slightly different than with the **Eager Recharging** strategy. This is due to the fact that there were no delays due to recharges early on, and thus, alternative routing choices were examined and a better route was eventually realized.

There are several reasons that make the **Lazy Recharging** strategy promising:

Allowing vehicles to exhaust their batteries enables the scheduling policy to consider alternative plans for a longer time. This maximizes the probability that client requests that would best be served together, i.e., nearby client requests, are actually assigned to the same vehicle. Moreover, visiting recharging stations only when it is absolutely necessary limits the number of total visits, as some vehicles may avoid the need for a recharge along their route. This limits the total route length, as each visit to a recharging station is associated with a distance cost. Exploiting the full battery capacity of each vehicle limits the probability that a second visit to a recharging station is needed for the same vehicle, which limits the total route length as well.

On the downside, the **Lazy Recharging** strategy does not consider proximity when scheduling visits to recharging stations. This can lead to routes that schedule visits to recharging stations after deliveries that are located faraway from any available recharging station. Therefore, vehicles may have to traverse long distances in order to restore the energy level of their batteries.

3.3 Smart Recharging

Our last strategy, termed **Smart Recharging**, is a combination of the **Eager** and **Lazy Recharging** strategies, aiming to capitalize on the advantages of both. In particular, this strategy attempts to exploit the capacity of the vehicle's battery and the proximity of delivery requests to recharging stations. To this end, the **Smart Recharging** strategy schedules proximity-driven visits to recharging stations. However, decisions based on proximity are only taken when a significant part of a vehicle's battery capacity is exhausted. That is, the **Smart Recharging** strategy examines the possibility of visiting a recharging station, even though the energy level of the vehicle's battery allows for subsequent requests, if the charging level is below 35%. As the case is with both **Eager** and **Lazy Recharging** strategies, a visit to a recharging station is scheduled when the battery is close to being exhausted with **Smart Recharging** as well.

Figure 5c shows a route formed after following the **Smart Recharging** strategy. We observe that this strategy led to a route that resembles the route formed with the **Lazy Recharging** strategy in its first part, and the route formed with the **Eager Recharging** strategy in its second part. In particular, the recharging station close to the first delivery location was ignored, as the energy level of the

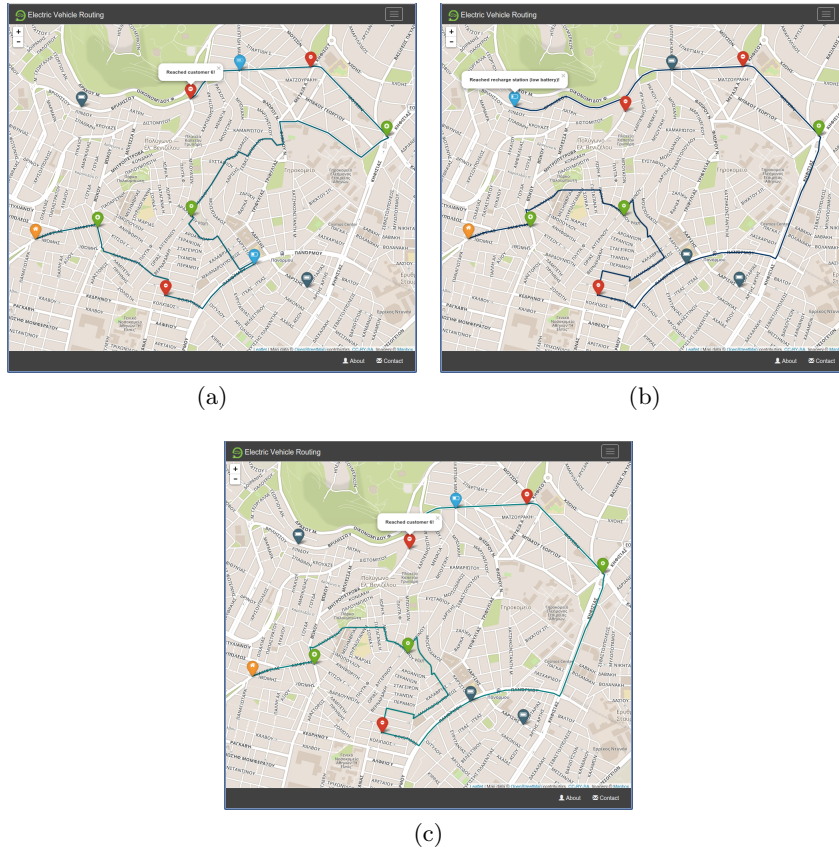


Fig. 5. Routing of a vehicle using **Eager** (5a), **Lazy** (5b), and **Smart Recharging** (5c) strategies, as depicted in our Web Application. **Smart Recharging** allows room for better routing decisions, while also favoring the avoidance of costly visits to faraway recharging stations that necessarily occur when the battery of a vehicle is exhausted.

vehicle was more than half full. This allowed for the better routing that occurred with **Lazy Recharging**. The recharging station close to the second delivery location is, however, visited, as the energy level at that stage is lower than 35%. This recharge restores the energy level of the vehicle, and a costly visit to a recharging station after the third delivery is avoided, as the case is with **Eager Recharging** as well.

4 Experimental Evaluation

We first present the experimental environment, the dataset and the settings that we applied to evaluate our on-line algorithm. Then, we proceed with the

evaluation of all our proposed recharging strategies by answering the following questions:

- How many vehicles does each strategy require for the instances of our dataset?
- How close to the solution that does not consider the need for visits to recharging stations is each one of our strategies with regard to the total mileage of the fleet?
- What is the average level of battery that each strategy opts to perform a visit to a recharging station at?
- How many visits to recharging stations does each strategy schedule?

4.1 Experimental Setting

We implemented and ran our on-line algorithm and the associated recharging strategies using Java 8. The experiments were carried out on a computer with an Intel[®] Core[™] i5-4590, with a CPU frequency of 3.30GHz, a 6MB L3 cache and a total of 16GB DDR3 1600MHz RAM and the Linux Xubuntu 14.04.03 (Trusty Tahr) x86 64 OS. The dataset that was used for the experiments comprises the 10-hour instances of Mitrović-Minić et al. [20] and is publicly available³. The set contains a total of 40 instances with 100, 300, 500, and 1,000 client requests. The service area is 60×60 km², with few delivery locations (around 6%) outside this area. The vehicle speed is 60 km/h. The vehicle fleet is empty at the beginning of the algorithm and vehicles are added as client requests arrive. The initial point of each vehicle (*depot*) is set to be (20,30), as in [20], to ensure that all transportation requested are serviced.

Given one of these instances, we determine the locations of seven recharging stations by placing one in the *depot* and the other six in two quadrants of the service area, in a random manner. The latter, serves the purpose of considering areas that do not provide access to recharging stations, which is expected in real-life situations. The maximum distance of a recharging station from a delivery location, at which the **Eager** and **Smart Recharging** strategies allow visits before battery exhaustion is set to 2.0 km. This value, was experimentally found to consistently allow more but not excessive recharges when compared to the **Lazy Recharging** strategy. For the parameters associated with the batteries of electric vehicles we adopt the criteria specified by Schneider et al. [29]. In particular, we set the battery capacity to the maximum of the following two values: (1) the charge needed to travel 60% of the average route length of the solution using our algorithm without energy constraints, and (2) twice the amount of battery charge needed to travel between a customer location and a station. Thus, we ensure that some vehicles will require a visit to a recharging station. Table 1 depicts the average values used for our dataset, expressed in terms of vehicle range (km). The consumption rate of the vehicles is set to 1.0. Finally, we consider batteries that recharge linearly with time and set the time needed

³ http://www.sfu.ca/~snezanam/personal/PDPTW/TestInstances/Rnd6_1h-2h-4h-6h-7h-Req/.

Table 1. Vehicle range and route length comparison.

No. of clients	Range	Unlimited Battery	Eager Recharging	Lazy Recharging	Smart Recharging
100	146.78	2,780.94	3,694.72	3,677.31	3,677.31
300	165.03	6,907.15	8,806.8	8,770.21	8,756.77
500	183.63	10,408.55	12,987.84	12,797.54	12,783.14
1,000	199.54	17,895.83	22,454.87	22,055.74	22,016.45

for a complete recharge to be equal to three times the average customer service time of the respective instance.

4.2 Fleet Size Comparison

We begin our evaluation by examining the impact of all proposed strategies on the minimization of the fleet size. We observe in Figure 6a that all three strategies which consider visits to recharging stations require significantly more vehicles to serve the customer requests than what would be needed if the battery capacity was not a concern. In particular, using vehicles that do not require recharging along their routes, we would need a fleet of approximately 45% less vehicles for the instances of all client request sizes.

Figure 6a also depicts that the **Eager Recharging** strategy requires the largest fleet in most cases. This indicates that attempting to reduce the fleet size by recharging sooner, and thus, performing more but smaller visits to recharging stations is not a good strategy in the long term. The **Lazy Recharging** strategy is more competitive, offering mild improvements over **Eager Recharging**. However, **Smart Recharging** proved to be the most successful strategy with regards to the fleet size, by outperforming **Lazy Recharging** in all cases, with the exception of the 100 client request instances, where there was a tie.

4.3 Route Length Comparison

We proceed by examining the impact of our strategies on the minimization of the route length. Table 1 shows the total mileage required for each of our three strategies, as well as the case of using battery with unlimited capacity, to service the client requests of our dataset. We observe that the increase inflicted upon the total route length due to visits to recharging stations is not as significant as the increase in the size of the fleet. This shows that all three strategies are still able to limit the total mileage spent; the time vehicles spend in recharging stations to restore their energy levels is more costly than the time needed to actually visit a recharging station.

The **Smart Recharging** strategy again stands out and outperforms both **Eager** and **Lazy** recharging. Therefore, we observe that **Smart Recharging** exhibits the best performance with respect to both the fleet size and the route length. This confirms our intuition that **Smart Recharging** benefits from the positive aspects of both **Eager** and **Lazy** recharging.

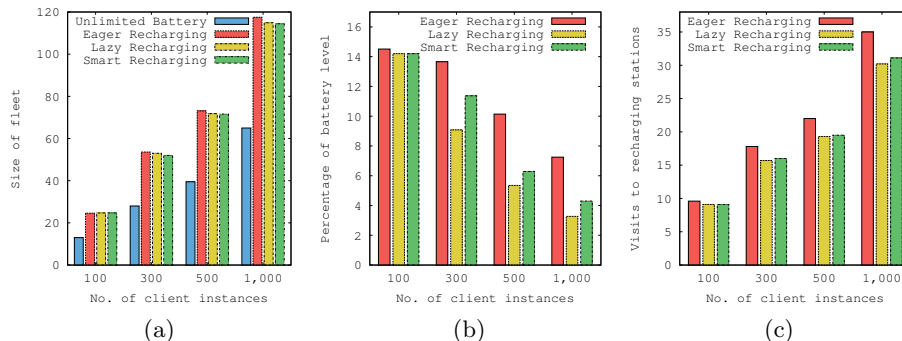


Fig. 6. Comparison of the fleet size needed (6a), the energy level of vehicles (6b), and the average number of recharges performed (6c) when using our three proposed strategies and a strategy without battery constraints. We observe that the impact of visiting recharging stations is significant, and that the **Smart Recharging** strategy stands out by outperforming **Eager** and **Lazy** recharging.

4.4 Battery Level & Total Recharges Comparison

The lifespan of batteries is susceptible to frequent discharges to deep levels [19]. Therefore, when maintaining an electric vehicle fleet it is important to adopt policies that slow down the reduction of usable battery capacity to extend the battery life and preserve the vehicle range. Figure 6b depicts the average percentage of the energy level of vehicles just before they visited a recharging station. The results regarding the **Lazy Recharging** strategy are alarming. There were cases when this strategy led to critically low energy levels (below 2%) before scheduling recharges. This indicates that although the **Lazy Recharging** strategy is very competitive, it may have a negative impact on the lifespan of the batteries of vehicles. The **Eager** and **Smart Recharging** strategies maintained a higher average energy level than **Lazy Recharging**. Therefore, they both may contribute to the reduction of the fleet’s maintenance costs in the long term.

For reasons of completeness, we also present in Figure 6c the total number of visits to recharging stations the vehicles performed when following each of the three strategies, for all the instances of our dataset. As was expected, we observe that the **Eager Recharging** strategy led to more visits than the **Smart Recharging** strategy, which in turn, led to slightly more visits than the **Lazy Recharging** strategy.

5 Related Work

Problems related to transportation, such as traffic congestion and air pollution, are increasingly troubling city authorities. This has led to the development of urban economy solutions that benefit on cloud-computing and the use of mobile devices to offer systems that minimize transportation cost [30,13].

The ever-increasing interest of companies in green logistics practices has driven carriers to the adoption of EVs. As typical delivery tours surpass the range these vehicles are able to provide, and their recharging time is considerable, scheduling policies must also handle visits to recharging stations. Gonçalves et al. [14] make a first attempt towards the investigation of the additional constraints imposed in a vehicle routing problem when BEVs are taken into account. While focusing on the PDPTW, [14] considers a mixed fleet of battery-powered EVs and conventional vehicles and assumes a limited driving range and a realistic charging time for vehicles. However, the proposed model allows BEVs to recharge anywhere, i.e., the locations of the recharging stations are not specified.

Erdoğan and Miller-Hooks [10] propose the *Green VRP (G-VRP)* that focuses on alternative fuel-powered vehicle fleets and study the effects of limited vehicle driving ranges in conjunction with limited refueling infrastructure. The G-VRP is formulated as a mixed integer linear program and features two construction heuristics: the *Modified Clarke & Wright Savings* heuristic as well as the *Density-Based* clustering algorithm working along with a customized improvement technique. However, as the G-VRP emphasizes on alternative fuel vehicles and not specifically on BEVs, the charging delays experienced by EVs are not considered.

Conrad and Figliozzi [5] introduced the Recharging Vehicle Routing Problem (RVRP). Vehicles have limited range and charging times are considered. However, recharging can only occur at certain customer locations. Schneider et al. [29] examine the electric vehicle routing problem with time windows and recharging stations. The assumptions are similar to those presented in [5]. However, the recharging stations of [29] are not located at customer locations. The proposed approach combines a variable search neighborhood algorithm with a tabu search heuristic. Regarding the recharging stations, [29] uses a new problem-specific neighborhood operator, called *stationInRe*, that performs insertions and removals of recharging stations.

To the best of our knowledge, this is the first work to consider dynamic pickups and deliveries that are serviced with the help of an EV fleet and whose vehicles require visits to recharging stations positioned at specific locations dispersed along a urban area. We address the problem by proposing three novel strategies, namely **Eager**, **Lazy**, and **Smart Recharging** for handling the vehicles while on traveling. We evaluate them experimentally and present the merits of each of the aforementioned choices.

6 Conclusion and Future Work

In this paper, we focus on urban distribution using electric vehicles and develop a system for the coordination and monitoring of the different parties involved in it. In compliance with introduced policies for green logistics and for economic reasons, transportation and logistics companies intend on deploying –or have already introduced– EVs in their pickup and delivery fleets. A vehicular fleet made up entirely of EVs adds several new dimensions to the classic vehicle rout-

ing problems. To this end, vehicles portray limited traveling range and call for power recharges along their routes to designated stations or service centers. The expected average charging time for replenishing the power supply to a vehicle may last for several hours. Also, recharging stations are far less common than gas stations and this is not expected to change soon. Therefore, routing and scheduling an EV fleet to satisfy dynamic transportation requests is far from trivial.

We build a system that enables communication with both customers and vehicles and propose three heuristic-based recharging strategies to tackle this problem. The first two, namely **Eager** and **Lazy Recharging** reflect two extreme situations: eagerly seeking nearby recharging stations to avoid lengthy visits to faraway stations, and exhausting the full battery capacity of vehicles before scheduling a recharge, respectively. Our third strategy termed **Smart Recharging**, combines advantages from both **Eager** and **Lazy Recharging**, and manages to outperform them in terms of both size of the fleet and total route length. We also study the impact that our proposed strategies have on the lifespan of batteries, and verify our intuition that **Lazy Recharging** may indeed speed up battery deterioration; the other two strategies are however likely to extend the lifetime of batteries.

We could further explore the potential of each of our three proposed strategies by additionally employing tabu-search heuristics [12]. In particular, we can periodically perform a tabu-search procedure after our cheapest insertion procedure, to examine which strategy benefits the most from reassigning requests to other vehicles. Furthermore, our algorithm does not currently consider deviations from customer time windows, which may occur in real-life situations. We can handle such cases by incorporating a penalty proportional to these deviations in our objective function.

Acknowledgments

This work has been partially supported by the EU H2020 “*GALENA*” grant with agreement n. 641515.

References

1. European Commission: Study on Urban Freight Transport. Technical report, April 2012.
2. London Freight Plan. Technical report, 2012.
3. J. Branke, M. Middendorf, G. Noeth, and M. Dessouky. Waiting Strategies for Dynamic Vehicle Routing. *Transportation Science*, 39(3):298–312, 2005.
4. M. Bruglieri, A. Colorni, and A. Luè. The Vehicle Relocation Problem for the One-way Electric Vehicle Sharing: an Application to the Milan Case. *Procedia-Social and Behavioral Sciences*, 111:18–27, 2014.
5. R. G. Conrad and M. A. Figliozzi. The Recharging Vehicle Routing Problem. In *Proceedings Industrial Engineering Research Conference*, 2011.

6. G. B. Dantzig and J. H. Ramser. The Truck Dispatching Problem. *Management Science*, 6:80–91, 1959.
7. R. Dekker, J. Bloemhof, and I. Mallidis. Operations Research for Green Logistics—An overview of Aspects, Issues, Contributions and Challenges. *European Journal of Operational Research*, 219(3):671–679, 2012.
8. G. Duleep, H. van Essen, B. Kampman, and M. Grünig. Impacts of Electric Vehicles, Deliverable 2, Assessment of Electric Vehicle and Battery Technology. *CE Delft*, 2011.
9. Y. Dumas, J. Desrosiers, and F. Soumis. The Pickup and Delivery Problem with Time Windows. *European Journal of Operational Research*, 54(1):7–22, 1991.
10. S. Erdoğan and E. Miller-Hooks. A Green Vehicle Routing Problem. *Transportation Research Part E: Logistics and Transportation Review*, 48(1):100–114, 2012.
11. W. Feng and M. 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.
12. M. Gendreau, A. Hertz, and G. Laporte. A Tabu Search Heuristic for the Vehicle Routing Problem. *Management science*, 40(10):1276–1290, 1994.
13. G. Gidófalvi, T. B. Pedersen, T. Risch, and E. Zeitler. Highly scalable trip grouping for large-scale collective transportation systems. In *EDBT 2008, 11th International Conference on Extending Database Technology, Nantes, France, March 25-29, 2008, Proceedings*, pages 678–689, 2008.
14. F. Gonçalves, S. R. Cardoso, S. Relvas, and Barbosa-Póvoa. Optimization of a Distribution Network using Electric Vehicles: A VRP problem. In *Proc. of 15 Congresso da Associação Portuguesa de Investigação Operacional*, Lisbon, Portugal, 2011.
15. O. H. Hannisdahl, H. V. Malvik, and G. B. Wensaas. The Future is Electric! The EV Revolution in Norway - Explanations and Lessons Learned. In *Proc. of the World Electric Vehicle Symposium and Exhibition (EVS27)*, pages 1–13, Barcelona, Spain, November 2013. IEEE.
16. International Post Corporation. Green Flash. <https://www.ipc.be/~media/Documents/PUBLIC/market-flash/Green%20Issues/GF18.pdf>. [accessed 22-December-2015].
17. P. R. Kleindorfer, K. Singhal, and L. N. Wassenhove. Sustainable Operations Management. *Production and Operations Management*, 14(4):482–492, 2005.
18. S. M. Lukic, J. Cao, R. C. Bansal, F. Rodriguez, and A. Emadi. Energy Storage Systems for Automotive Applications. *IEEE Transactions on Industrial Electronics*, 55(6):2258–2267, 2008.
19. A. Millner. Modeling Lithium Ion Battery Degradation in Electric Vehicles. In *Proc. of the 2010 IEEE Int. Conf. on Innovative Technologies for an Efficient and Reliable Electricity Supply (CITRES'10)*, pages 349–356, Waltham, MA, 2010. IEEE.
20. S. Mitrović-Minić, R. Krishnamurti, and G. Laporte. Double-Horizon Based Heuristics for the Dynamic Pickup and Delivery Problem with Time Windows. *Transportation Research Part B: Methodological*, 38(8):669–685, 2004.
21. S. Mitrović-Minić and G. Laporte. Waiting Strategies for the Dynamic Pickup and Delivery Problem with Time Windows. *Transportation Research Part B: Methodological*, 38(7):635–655, 2004.
22. S. Pelletier, O. Jabali, and G. Laporte. Goods Distribution with Electric Vehicles: Review and Research Perspectives. Technical report, CIRRELT: Interuniversity

- Research Centre for Enterprise Networks, Logistics and Transportation, Montreal, Canada, September 2014.
23. B. G. Pollet, I. Staffell, and J. L. Shang. Current Status of Hybrid, Battery and Fuel Cell Electric Vehicles: from Electrochemistry to Market Prospects. *Electrochimica Acta*, 84:235–249, 2012.
 24. H. N. Psaraftis. Dynamic Vehicle Routing Problems. In *Vehicle Routing: Methods and Studies*, chapter 16, pages 223–248. Elsevier Science Publisher B.V. (North Holland), 1988.
 25. H. N. Psaraftis. Dynamic Vehicle Routing: Status and Prospects. *Annals of Operations Research*, 61(1):143–164, 1995.
 26. M. Savelsbergh and M. Sol. DRIVE: Dynamic Routing of Independent Vehicles. *Operations Research*, 46(4):474–490, 1998.
 27. M. W. Savelsbergh. Local Search in Routing Problems with Time Windows. *Annals of Operations research*, 4(1):285–305, 1985.
 28. A. Sbihi and R. W. Eglese. Combinatorial Optimization and Green Logistics. *Annals of Operations Research*, 175(1):159–175, 2010.
 29. M. Schneider, A. Stenger, and D. Goeke. The Electric Vehicle-Routing Problem with Time Windows and Recharging Stations. *Transportation Science*, 48(4):500–520, 2014.
 30. O. W. Shuo Ma, Yu Zheng. T-share: A large-scale dynamic taxi ridesharing service. ICDE 2013, April 2013.
 31. T. M. Sweda, I. S. Dolinskaya, and D. Klabjan. Optimal Recharging Policies for Electric Vehicles. Technical report, Northwestern University, Dept. of Industrial Engineering and Management Sciences, Evanston, IL, November 2014.
 32. UberPool. Announcing UberPool. <https://newsroom.uber.com/announcing-uberpool/>. [accessed 22-December-2015].
 33. U.S. Department of Energy. Plug-in Electric Vehicle Handbook for Fleet Managers. http://www.afdc.energy.gov/pdfs/pev_handbook.pdf. [accessed 22-December-2015].