REC: Predictable Charging Scheduling for Electric Taxi Fleets

Zheng Dong[†], Cong Liu[†], Yanhua Li[‡], Jie Bao[§], Yu Gu[¶], Tian He[*]
[†]The University of Texas at Dallas, [‡]Worcester Polytechnic Institute
[§]Microsft Reseach Asia, [¶]Visa USA, [*]University of Minnesota

Abstract. Due to the energy security concern, our society is witnessing a surge of EV fleet applications, e.g., public EV taxi fleet systems. A major issue impeding an even more widespread adoption of EVs is range anxiety, which is due to several factors including limited battery capacity, limited availability of battery charging stations, and long charging time compared to traditional gasoline vehicles. By analyzing our accessible real-world EV taxi system-wide datasets, we observe that current EV taxi drivers often suffer from unpredictable, long waiting times at charging stations, due to temporally and spatially unbalanced utilization among charging stations. This is mainly because current taxi fleet management system simply rely on taxi drivers to make charging decisions. In this paper, In this paper, we develop REC, a Real-time Ev Charging scheduling framework for EV taxi fleets, which informs each EV taxi driver at runtime when and where to charge the battery. REC is able to analytically guarantee predictable and tightly bounded waiting times for all EVs in the fleet and temporally/spatially balanced utilization among charging stations, if each driver follows the charging decision made by REC. Moreover, REC can further efficiently handle real-life issues, e.g., allowing a taxi driver to charge at its preferred charging station while still guaranteeing balanced charging station utilization. We have extensively evaluated REC using our accessible real-world EV taxi system-wide datasets. Experimental results show that REC is able to address the unpredictability and unbalancing issues existing in current EV taxi fleet systems, yielding predictable and tightly bounded waiting times, and equally important, temporally/spatially balanced charging station utilization.

1 Introduction

Energy security is one of the biggest issues in the global political climate. Instability in global oil producing nations has driven the need for major energy importing nations to be less reliant on foreign sources of energy. The transportation sector, which currently accounts for nearly 72 percent of global oil demand [17], is one such sector that has seen a major impetus to transform. Many countries have been promoting the usage of electric vehicles (EVs). For instance, U.S. is expected to have one million EVs by 2015 [19]. More recently, our society is witnessing a surge of EV fleet applications, where a fleet of EVs is managed and operated as a single entity. A particular example is the public EV taxi fleet system (e.g., the Electric Cab Corporation), where EV taxis are used to replace traditional gasoline taxis for serving customers.

EVs have the potential to alleviate the energy security concerns and reduce city pollution by having zero tail pipe emissions. Most EVs nowadays are equipped with multiple lithium-ion battery modules as the power source [18]. Currently, there exist a major issue impeding the widespread adoption of EVs, i.e., range anxiety, which is largely due to the combined impact of the limited battery capacity and the limited availability of battery charging stations in most cities. Due to range anxiety, existing EV taxi fleet management systems simply allow taxi drivers to freely make charging decisions, i.e., when and where to charge the battery. Unfortunately, such unpredictable and often greedy charging decisions made by individual drivers may easily cause certain charging stations over-utilized, while left some other stations under-utilized.

Key motivation of this work. We are fortunate to have access to real-world EV taxi system-wide datasets, which include EV taxi trajectory data, road map data, and charging station data. These datasets were collected from the same time window from November 1^{st} - 30^{th}, 2013, in Shenzhen City, China. The datasets were generated from 550 EV taxis equipped with GPS sets consisting 23,967,501 GPS records. By thoroughly analyzing these datasets (data analysis details are given in Sec. 2), we conclude that current EV taxi drivers often suffer from unpredictable, often long waiting times at charging stations, due to the unbalanced charging station utilization both spatially and temporally. This undesirable consequence is intuitive because certain charging stations and certain time periods are always preferred by many taxi drivers due to common sense principles, e.g., most drivers may have the same tendency of driving to a large charging station with more charging poles. This unpredictably long waiting time issue is significantly exaggerated by the fact that compared to gasoline vehicles, it takes a much longer time to fully charge an EV under current charging techniques. For instance, even with the latest supercharger technology, it still takes 30-75 minutes to fully charge an EV battery [23].

Making individual charging decisions may not be an issue for traditional gasoline EV taxis, because there is a sufficient number of gasoline stations in most cities. For EV taxis, unfortunately, the limited number of charging stations existing in most cities make the charging issue rather challenging. Clearly, simply replying on individual taxi drivers to make charging decisions will result in bad performance in terms of unpredictable waiting time at charging stations, and city-wide unbalanced charging station utilization. This will further cause both individual taxi drivers and the fleet system lost revenue. To resolve this challenge, this paper investigates the following question: how to design an EV charging scheduling protocol that guarantees predictable and tightly bounded waiting time for all EV taxis in a taxi fleet system. Using a heuristic-based scheduler is undesirable because it cannot guarantee predictable system performance: individual taxi drivers do not know the waiting time at a charging station scheduled by the system, thus may refuse to follow the charging scheduling decisions given to them. Moreover, guaranteeing predictable waiting times at charging stations may often imply a sufficiently balanced utilization among all available charging stations, particularly given a limited number of charging stations. Thus, designing a fleet-wise charging scheduling solution that guarantees predictable and tightly bounded waiting times for all EVs in the system is critical in making any such solutions be practically implementable in real-world scenarios.

In this paper, we develop REC, a Real-time Ev Charging scheduling framework for an EV taxi fleet system, which informs each EV taxi driver at runtime when and where to charge the battery. REC is able to guarantee predictable and tightly bounded waiting times for all EVs in the fleet if each driver follows the charging decision made by REC. To reach this goal, REC employs a fast charging scheduling policy, fundamentally motivated by the existing real-time scheduling theory on scheduling a set of hard real-time suspending tasks on multiprocessors, due to the equivalence of our charging scheduling problem and this classical real-time CPU scheduling problem. Moreover, REC can further efficiently handle real-life issues, e.g., what if a taxi driver has a strong preference on a specific charging location and refuses to charge at the REC-scheduled location? REC further develops a novel taxi migration plan that allows drivers to charge at their preferred charging station, while still guaranteeing balanced station utilization through migrating a minimum set of EVs among charging stations. To compensate drivers who are chosen to migrate to a new charging station, REC identifies the best migration plan for each such driver in terms of various compensation objectives, e.g., a minimum migration-induced detour distance or migrating to a hot-spot charging location with busy trip demands predicted using historical data.

Our specific contributions are listed as follows:

- To the best of our knowledge, this is the first solution studying the charging scheduling problem of a fleet of EVs, which yields predictable waiting times at charging stations for all EVs in the system.
- Our developed charging scheduling protocol REC is rather simple, thus being practically implementable in real-world taxi fleet systems. Despite its simplicity, it analytically guarantees tightly bounded waiting times at charging stations and is capable of handling real-life issues including giving the option to certain taxi drivers to choose their own preferred charging locations while still guaranteeing predictable waiting times.
- We have extensively evaluated REC using our accessible real-world EV taxi system-wide datasets. Experimental results show that REC is able to address the unpredictability and unbalancing issues found through analyzing the original datasets, yielding very predictable and tightly bounded waiting times, and equally important, both spatially and temporally balanced charging station utilization.

2 Motivation

Traditional gasoline taxi systems provide a primary transportation service in modern cities. Most street taxis respond to passengers' requests on their paths, and take passengers to their specified destinations. Therefore, traditional taxi systems rely

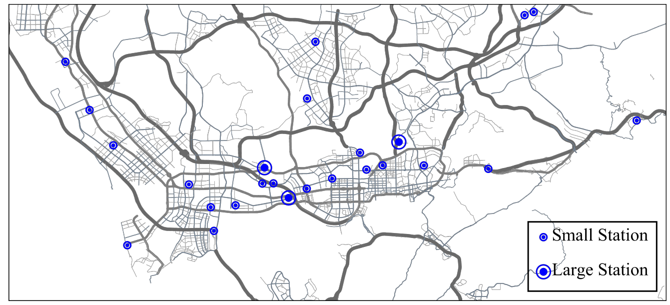

Fig. 1: The road map of Shenzhen city.

on drivers to drive around and arbitrarily pick up passengers on streets. This service model has successfully served up to 25% public passengers in metropolitan areas, such as San Francisco and New York [16]. In a traditional taxi system, a couple of key dynamics affect their service quality: a) dynamic passenger demand. The spatiotemporal patterns of demand include both regular factors, such as rush hours and busy areas, and irregular ones, such as weather, traffic, holiday schedule, etc. b) dynamic taxi supply. Taxis have different mobility patterns, since drivers have different working schedules.

Comparing to traditional taxis, the long recharging time of EV leads to the forming of a gap between EV and its supplying equipment. Thus, the electric taxi system introduces a third factor on their service quality: when and where to re-charge the battery. Intuitively, if drivers follow their individual schedule and choose the popular charging stations for recharging, they often end up queueing in the charging stations and suffer unpredictable and long waiting time. This will further cause both temporally- and spatially-unbalanced demands among charging stations.

We now describe our accessible real-world EV taxi system-wide datasets, and present the potential issues and insights observed from analyzing the datasets. Then, we formally formulate the problem and outline our solution framework.

2.1 Data Description

We have three accessible sets of data for the analysis, including (1) EV taxi trajectory data, (2) road map data, and (3) existing charging station data. All these datasets were collected from the same time window, i.e., November 1^{st} -30^{th}, 2013. Below, we describe each of the datasets in detail.

EV taxi trajectory data. The EV taxi trajectory dataset consists of sequential GPS records for electric taxis, which was collected during November 1st-30th, 2013 in Shenzhen City, China. The dataset was generated from 550 electric taxis equipped with GPS sets consisting 23,967,501 GPS records of EV taxis. For EV, the average GPS recording period is about 40 seconds. Each GPS record contains five useful fields for our study, including the taxi ID, time stamp, latitude, longitude, passenger load indicator (PLI), where PLI is a binary variable indicating whether or not a taxi is taking passenger(s).

Road map data. We obtain a bounding box of Shenzhen city through Google Geocoding API [1]. The south-west and north-east corners of the bounding box are (22.447203, 113.748964) and (22.83385, 114.601127) in latitude and longitude covering an area of roughly 1,804 km^2. Fig. 1 shows the roads within

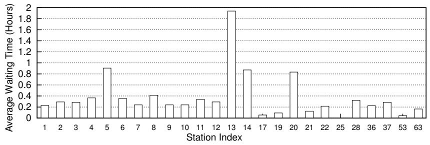

(a) The average waiting time between 0:00 to 4:00.

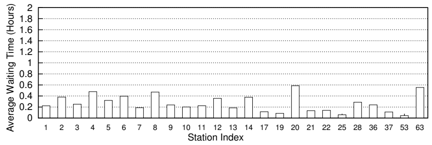

(b) The average waiting time between 4:00 to 8:00.

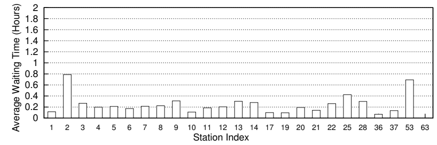

(c) The average waiting time between 8:00 to 12:00.

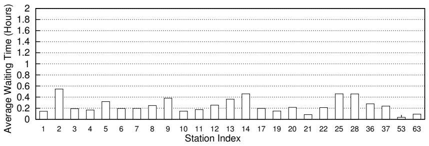

(d) The average waiting time between 12:00 to 16:00.

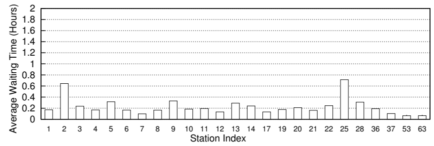

(e) The average waiting time between 16:00 to 20:00.

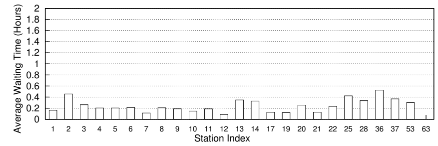

(f) The average waiting time between 20:00 to 24:00.

Fig. 2: The average waiting time at different charging stations during different time periods.

the bounding box. With the bounding box, Shenzhen road map data were obtained from OpenStreetMap [2], which contains all road segments and their road types.

Charging station data. Within the bounding box, there are 25 charging stations deployed and in use throughout November 2013. The spatial distribution of these charging stations is presented in Fig. 1, in which the number of deployed charging points is indicated with different marker size. The large circle represents the charging stations with more than 50 charging poles and the small ones indicate the charging stations with less than 50 charging poles. As we could observe in Fig. 1, there are three large charging stations and 22 small charging stations.

2.2 Data Analysis and Observed Issues

For EV taxi drivers, one of the most critical issue they concern about is the amount of waiting time wasted at charging stations. Thus, we first analyzed these three datasets to check whether the waiting time for each EV at charging stations is reasonable. We them conducted further data analysis to identify potential issues of the charging decisions made by taxi drivers.

To conduct such analysis, we need to first verify whether an EV taxi is inside a charging station during each time slot. This can be simply analyzed by examining whether the geolocation of a trajectory is inside the range of any charging station. Secondly, we need to determine whether an EV taxi is charging or waiting within a charging station. We assume that for each EV taxi, the last longest idle period within a charging station area is considered to be the charging period while other idle periods are considered to be waiting periods. We are thus able to extract the average waiting time across different charging stations during different time intervals. The results are shown in Fig 2. We observe that the average waiting time varies significantly during different time intervals and among different charging stations. In many cases, the average waiting time can be viewed as long. For example, as seen in Fig. 2a, the average waiting time at station 13 is almost 2 hours, while the waiting time at station 25 is almost 0.

To understand the fundamental reasons behind these observations, we further summarize the results of Fig. 2 and present them separately in Fig. 3 and Fig 4, which reveal two major charging scheduling issues for EV taxis, including (i) spatially unbalanced charging station utilization, (ii) temporally unbalanced charging station utilization.

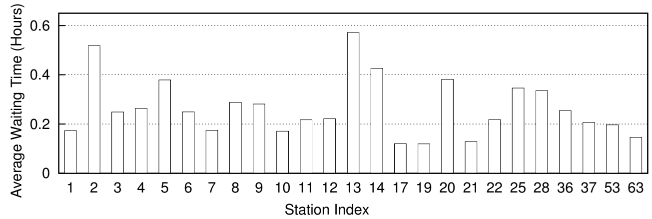

Fig. 3: Average waiting time when charging at different charging stations, demonstrating significant spatial imbalance of EV charging loads.

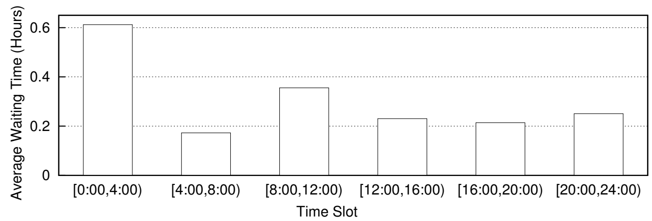

Fig. 4: Average waiting time when charging at different time of the day, demonstrating significant temporal imbalance of EV charging loads.

Spatially unbalanced charging station utilization. As shown in Fig. 3, the average waiting time of EV taxis varies significantly at different charging stations. For instance, the average waiting times at stations 2, 5, 13, 14 are considerably longer than many other stations. This implies that more EV taxis drivers tend to charge batteries at these stations, thus causing over-utilization. By analyzing the data, the potential reasons behind are due to 1) the scale of a charging station, 2) the chance of serving more customers if charging at a certain charging station, and 3) the distance between taxi location and a charging station. In many cases, a taxi driver tends to make greedy charging decisions, e.g., drive to nearest large-scale charging station. Such greedy choices unfortunately often cause decreased and unpredictable system-wide performance.

Temporally unbalanced charging station utilization. Moreover, Fig. 4 shows the average waiting time among all charging stations over six time slots (each covers four hours of a day). As seen, the average waiting time also varies significantly over different time intervals within a day. For example, the average waiting time for EVs charged during time slots $[0:00, 4:00)$ and $[8:00, 12:00)$ is considerably longer compared to other time slots. A potential reason for this observation is because certain charging stations are reserved for electric bus/shuttle charging during those early morning time slots.

Insights obtained from the data analysis. The above two shortcomings observed from analyzing the EV datasets highlight the importance of developing a smarter EV fleet charging scheduling system. Balancing spatio-temporal taxi supply across the entire city is a design requirement. For each individual driver, unbalanced charging station utilization often yields unpredictable, long waiting time in charging stations, which results in a direct cost on the taxis' service time and thus revenue. To resolve these issues, we seek to develop a fleet-wide EV taxi charging scheduling system that recommends each taxi driver the time and location for recharging its EV battery. Our goal is to guarantee predictable and bounded waiting time at the charging station for each taxi in the system, and balance the charging demand both temporally and spatially.

2.3 Formal System Model on EV Charging Scheduling

We now mathematically describe our system model. We consider an EV taxi fleet with a set $S = T_1, T_2, T_3, ..., T_n$ of n EV taxis providing public transportation services in a city 24-hours a day. Each taxi has two operation modes: driving in the city to deliver passengers to their destinations and staying in the charging station for recharge. If the fleet management system simply relies on individual drivers to decide when and where to recharge batteries, the waiting time for EVs vary in a large range, e.g., from 0.1-hour to 1.95-hour according to our analyzed data shown in Sec. 2.2. A 1.95-hour waiting time is clearly unacceptable for any EV taxi driver. The fundamental reason why such unpredictable waiting time occurs is the utilization of charging stations suffers from spatial and temporal imbalance as we observed from the data sets. The over-utilized charging stations will cause increased waiting time at those stations, while the resources in other under-utilized charging stations get wasted. We thus propose to develop a fleet-wise real-time EV charging scheduling framework (REC), which schedules each EV taxi in the system to one of the m charging stations spatially distributed in the city. Following the charging scheduling decisions made by REC, all EVs in the system will have predictable and tightly bounded waiting time at charging stations, which will further ensure a spatially and temporally balanced utilization among all charging stations. To be practically implementable in real-world, REC is designed to be capable of handling real-life issues including giving the option to certain taxi drivers to choose their own preferred charging locations while still guaranteeing predictable waiting times.

2.4 Intuitive Ideas behind REC

Before describing our detailed design of REC, we first explain a key intuition behind REC. The EV charging scheduling problem becomes rather trivial to resolve if we make the following hypothesis: EVs can be instantly transported to a charging station with no cost. Then, a simple yet effective solution would be to instantly transport an EV when its battery becomes empty to a charging pole with the shortest waiting time. This will yield the charging demand among charging stations temporally and spatially balanced, while guaranteeing a minimum waiting time for each EV in the system. This hypothesis borrows the concept of "work-conserving" scheduling from multiprocessor real-time scheduling. A multiprocessor scheduling algorithm is said to be work-conserving if it does not idle any processor when one or more jobs are waiting. Any work-conserving scheduling algorithm is known to be able to achieve good balanced resource utilization among processors [8]. Similarly, if REC can instantly transport each EV to a charging station with the smallest waiting time, then it is work-conserving and shall achieve the demand balancing goal. Thus, if this hypothesis were realistic, we can simply dispatch the EVs under any work-conserving multiprocessor real-time scheduler and derive a predictable waiting time for each EV accordingly.

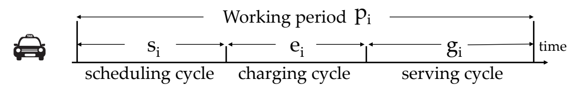

Fig. 5: Taxi working period.

Unfortunately, this hypothesis is clearly unrealistic for the EV scenario. It takes both time and energy for an EV to reach a charging station. Nonetheless, this hypothesis reveals a key challenge of the EV charging scheduling problem and an interesting idea. The challenge is that how to design a charging scheduling algorithm that achieves the predictability/balancing goal while considering the scheduling overhead, i.e., the time and energy used by an EV to reach its scheduled charging station. Our high-level idea is to leverage a new set of real-time scheduling techniques that allow tasks to suspend during their execution on processors. Thus, we can model each EV as a suspending task, where a suspension phase is used to model an EV's status when it travels to its scheduled charging station and an execution phase is used to model an EV's status when it is charged at the charging station.

Following this idea, we design REC, which is a two-layer scheduling framework, including a basic design layer that addresses the fundamental predictability/balancing issue, and an advanced design layer that considers real-world factors such as taxi drivers' preferences.

First layer of REC. The first layer of REC is an EV taxi charging scheduling algorithm, which informs each EV in the system its scheduled charging station when the EV's battery energy drops below a specific threshold. If all EVs in the system follow the scheduling decisions made by REC, then we will show that REC can guarantee predictable and tightly bounded waiting time for each EV in the system, and thus balanced charging demand among charging stations.

Second layer of REC. Although the first layer addresses the predictability/balancing issue, it assumes that taxi drivers will all obey the scheduling decisions made by REC. However, in real-world scenarios, drivers may often have their own preferences about where to charge their EVs. To address this important practical issue, the second layer of REC develops a migration technique, which allows any taxi driver of EV T_i to submit a preferred charging station at the time when the battery energy drops below the threshold. REC will then re-schedule and migrate EVs among charging stations accordingly to satisfy T_i's request while still guaranteeing predictable waiting times at charging stations for all EVs in the system. In such scenarios, if REC needs to migrate some

EVs to other charging stations in order to guarantee fleet-wise predictability, then REC will always find a migration plan that compensates such migrated EVs. We define a general compensation function that can be used to describe different optimization objects. We illustrate how to use this function to describe two example optimization objectives: 1) minimum migration overhead, which represents the additional time a migrated EV needs to reach its newly scheduled charging station, and 2) maximum short-term profit, which represents the time taken for a migrated EV to serve a passenger after charging at its re-scheduled charging station.

3 First Layer of REC

In this section, we present a charging scheduling algorithm that guarantees predictable waiting times for all EVs in the system. Our overall approach is to transform this problem into a real-time multiprocessor suspending task scheduling problem, and leverage existing solutions designed for the suspending task scheduling problem. We first introduce some concepts and terms that are needed to define this transformation.

3.1 Taxi working period

Before presenting our designed charging scheduling algorithm, we first describe when REC will inform a taxi where to charge its battery. During a taxi's working period, the driver may first serve various passengers. When its EV battery energy drops to a certain threshold (defined later), REC will advice the driver to drive to a charging station. The time taken by the driver to reach the scheduled charging station is defined to be a *scheduling cycle*, as illustrated in Fig. 5. After arriving at the charging station, the taxi may experience certain waiting time first and then start charging. After being fully charged, the taxi will begin serving passengers again. Since this pattern will repeat for each taxi, we define a working period of a taxi to be composed of a scheduling cycle, following by a waiting time cycle, then a charging cycle, and finally a serving cycle. Such a working period will repeat for each taxi. Note that different working periods may have different lengths. We now formally define these components.

Definition 1. *Let p_i denote T_i's working period, which is composed by a scheduling cycle with length s_i, a charging cycle with length e_i, and a serving cycle with length g_i. T_i experiences such working cycles in a repeating manner.*

Definition 2. *Based on the definitions of p_i and e_i for taxi T_i, the charging station utilization of T_i is defined as $u_i = \frac{e_i}{p_i}$ (note that $u_i < 1$). The total charging station utilization of all taxis in the system is given by $U = \sum_{i=1}^{n} \frac{e_i}{p_i}$. Intuitively, u_i characterizes the amount of charging resource T_i needs during each of its working periods.*

We now define the upper bounds of each component of a working period for later analysis purposes. Note that we will not upper bound the waiting cycle herein since we will show how to bound it under our developed charging scheduling algorithm.

Scheduling cycle s_i: In order to achieve balanced utilization of the charging stations, theoretically, an EV taxi shall be able to be dispatched to any charging station in the city. Thus, any scheduling cycle of T_i can be safely upper-bounded by the amount of time needed for T_i to travel between the two farthest locations on the city map, as illustrated by the red line shown in Fig. 6. According to our data sets, an EV can travel over 200 miles after a single charge. Since an EV can reach any charging station within 30 miles from anywhere in the Shenzhen City, an EV can start driving to its scheduled charging station when the battery has around 15% remaining energy (this threshold can be flexibly set to a larger number for safety purposes).

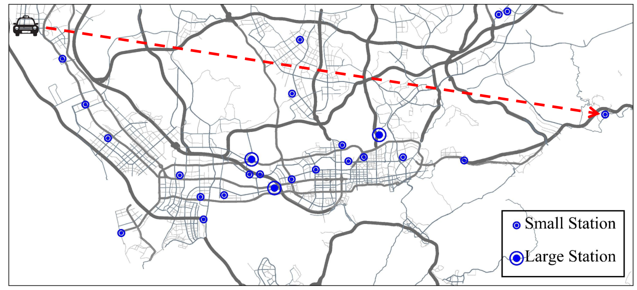

Fig. 6: Example illustrating the scheduling period.

Charging cycle e_i: When a taxi enters into a charging station, its battery is assumed to be at some low level. Thus, the charging cycle e_i of T_i is upper-bounded by the amount of time T_i takes to charge its battery from an empty status to a full status, which takes around 30-75 minutes under the latest supercharging techniques [3], [23].

Serving cycle g_i: According to the discussion on the scheduling cycle, EVs may consume around 85% (specific to the ShenZhen City environment) of the battery energy in during the serving cycles. The corresponding travel distance is at least 170 miles. Through analyzing our data sets, the serving cycle can thus be upper bounded by the amount of time used to drive 170 miles. (Note again that parameters including the serving cycle can be upper bounded under different scenarios following the same reasoning.)

Working period p_i: Given the definitions of the scheduling cycle s_i, charging cycle e_i, and serving cycle g_i, the working period p_i of a taxi T_i can be safely bounded by $s_i + e_i + g_i$.

3.2 REC's Charging Scheduling Algorithm

Without a judicious charging scheduling algorithm, multiple EVs may arrive at one charging station roughly at the same time, which will cause rather long waiting time for later-arrived EVs. More importantly, such waiting time is completely unpredictable, which may cause the overall fleet management to be orderless. Thus, we propose a smart charging scheduling algorithm to properly prioritize and schedule EV taxis to charging stations in real time, which guarantees each EV in the system a predictable and tightly bounded waiting time at any scheduled charging station. Besides this essential goal, our charging scheduling algorithm should also be simple enough to implement in practice. Particularly given a large number of EV taxis traveling in the city, any complicated charging scheduling algorithm may cause significant amount of runtime overheads to make real-time scheduling decisions.

We now present REC's charging scheduling algorithm, which is developed through transforming the charging scheduling problem into a real-time multiprocessor suspending task scheduling problem. We first introduce the real-time multiprocessing suspending task scheduling problem and then show the transformation.

Real-time multiprocessor suspending task scheduling. This CPU scheduling problem is to schedule a set of real-time sporadic tasks that may self-suspend (e.g., due to accessing shared resources) on a multiprocessor platform. Each sporadic task is released repeatedly, with each such invocation called a job. Jobs alternate between computation and suspension phases in any arbitrary manner. Each job of τ_i executes for at most e_i time units (across all of its execution phases) and suspends for at most s_i time units (across all of its suspension phases). The j^{th} job of τ_i, denoted $\tau_{i,j}$, is released at time $r_{i,j}$ and has a deadline at time $d_{i,j}$. Associated with each task τ_i are a period p_i, which specifies the minimum time between two consecutive job releases of τ_i, and a deadline d_i, which specifies the relative deadline of each such job, i.e., $d_{i,j} = r_{i,j} + d_i$. The utilization of a task τ_i is defined as e_i/p_i, which characterizes the long-term processor capacity requested by τ_i (i.e., a task with a utilization of 0.5 is expected to utilize half of the processor capacity over the long term).

Similarity between the two problems. The charging scheduling model can be equivalently viewed as the real-time suspending task scheduling model. Each charging pole can be viewed as a CPU core. Each EV can be viewed as a suspending task, where an EV's working period is equivalent to a suspending task's period. Within each working period of an EV, it incurs a scheduling cycle, a charging cycle, and a serving cycle, which correspond to a suspension phase, an execution phase, and another suspension phase of a suspending task, respectively. The charging activity can be equivalently viewed as task execution on a CPU core. Each working period of an EV can be equivalently viewed as a newly release job of a suspending task. The start time of each working period (also the start time of the scheduling cycle) can be viewed as a new job release of a suspending task.

An EV may experience waiting time after arriving at a charging station because all charging poles may be in use by other EVs. This corresponds to the scenario where a real-time task is interfered and delayed by other tasks that occupy all available CPU cores. Thus, in order to bound the wait time under the charging scheduling problem, we can alternatively seek for a scheduling algorithm that may bound the response time (thus the interference delay) of any job released by a suspending task under the real-time task scheduling scenario.

The global-earliest-deadline-first (GEDF) scheduler. GEDF is designed to schedule real-time suspending tasks on a multiprocessor [13]. Under NPEDF, a job $\tau_{i,j}$ with the smallest $d_{i,j}$ has the highest priority. Ties are broken by task ID. At each time instance, GEDF schedules the M jobs with the highest priorities among all release jobs on M CPU cores for execution.

Motivated by the above problem transformation, we apply GEDF in our problem context to prioritize EVs. Since EVs do not have deadlines, we set a deadline of the j^{th} working period (or the j^{th} job) of EV T_i to be the time when a working period ends (i.e., the end of the corresponding serving cycle). Whenever a new job is released (i.e., a new working period of an EV starts), REC will schedule it to a charging station with a charging pole that has the shortest waiting time. Note that this waiting time accounts for the charging cycle of any other EV that has already been scheduled to this charging pole, even if such EVs have not physically arrived at the charging station.

The rest of this section focuses on showing that GEDF is able to yield a waiting time for each EV tightly bounded by $p_i - e_i - s_i$. We recently introduced new scheduling analysis techniques showing that (Theorem 1 in [4]), any HRT suspending job $\tau_{i,j}$ can complete by the end of its period scheduled under GEDF, provided that $U_{sum} + \sum_{j=1}^{K} v^j \leq m - (m-1) \cdot \max_{\tau_i \in \tau}(\frac{e_i + s_i}{p_i})$ and $\exists K \ (2 \leq K \leq n), \sum_{i=1}^{\lfloor \frac{K}{2} \rfloor} E^i \geq p_{max}$, where U_{sum} is the total system utilization, v^j is the j^{th} maximum suspension ratio, where a task's suspension ratio is defined to be the ratio of its suspension length divided by its period, e_i, s_i, and p_i denotes task τ_i's execution time, suspension length, and period, respectively, p_{max} denotes the maximum task period in τ, and E^i denotes the i^{th} minimum $(e_j + s_j)$ among tasks $\tau_j \in \tau$.

As discussed above, the EV charging scheduling model can be equivalently viewed as the HRT suspending task model, each EV is guaranteed to complete both its scheduling phase and execution phase by the end of its period according to Theorem 1 in [4]. Thus, the waiting time for an EV T_i can be tightly upper bounded by $p_i - e_i - s_i$. In other words, if we use GEDF to schedule EVs to charging stations as explained above, then each EV is guaranteed to wait at its scheduled charging station for at most $p_i - e_i - s_i$ time units (recall that p_i, e_i and s_i are defined in Sec. 3.1).

Addressing the practical non-preemptive charging issue. GEDF is a preemptive scheduler, where a higher-priority task may preempt the execution of a lower priority task. For our EV charging scheduling problem, the above predictable waiting time bound can be guaranteed if any EV with higher priority arrives at its scheduled charging station no later than any other EV with a lower priority. If not, GEDF will allow the later-arrived higher priority EV to preempt another EV with a lower priority that is charging. Clearly, this behavior conflicts with the practice, where EVs are charged in a first-in-first-out order. It is not practical to assume that drivers may be willing to wait while there are charging poles available.

Thus, we further twist the above-mentioned analysis on bounding the waiting time to accommodate this practical consideration. Our key observation is the following: *any high-priority EV that has arrived at its scheduled charging station may be blocking by a low-priority EV for at most once*. This is because any such blocked high-priority EV T_i will get charged whenever a lower-priority EV finishes charging, unless there are other EVs with higher priorities than T_i arrive at the same charging station as T_i before T_i is assigned a charging pole. The later case is not a problem because GEDF intends to let this higher-priority EV to be charged earlier than T_i anyway. Due to this observation, we know that when this blocking (due to non-preemptive charging process) behavior is allowed in practice, the waiting time of any EV can still be bounded by $p_i - e_i - s_i + e_{max}$, where e_{max} represents the maximum charging time of any EV in the system. This additional term e_{max} is exactly due to the fact that any EV may be blocked by at most one lower-priority EV.

Despite the above-derived waiting time bounds, we believe a fundamentally more important goal achieved by this transformation-based solution is that it allows us to develop an efficient scheduling policy that has strong analytical properties in theory, thus having a higher possibility to execute well in practice compared to other pure heuristic-based approaches built on a rather weak foundation (as also proved by our extensive experimental results discussed in Sec. 5).

4 The Second Layer of REC

In Sec. 3, we have developed a transformation-based approach that can tightly bound the waiting time for any EV in the system, *provided that all EV drivers are willing to follow the scheduling decision given by our system*. However, in practice, a common observation is that different EV drivers may have their own preference on the charging location, and such preferences may be totally unpredictable. For example, if an EV driver has a very important appointment close to charging station A, but it is scheduled to be charged at another station B by REC, then this driver may be willing to suffer a longer waiting time at A and still choose to charge at A. The second layer of REC is designed to fully consider such real-life issues. Note that the requests can only be sent to REC at the time when an EV's energy level drops to its charging threshold.

REC is able to accommodate a EV driver's specific charging location request while the utilization of all charging stations remains spatially and temporally balanced. REC enables this capacity through incorporating a novel EV migration scheme, under which REC will change the charging location for a minimum set of EVs whenever an EV driver in the system requests a different charging location which is different from the scheduling decision made by REC. Since such migrations may negatively impact some EVs in the system (e.g., they may be de-toured and need to spend additional time to reach the newly assigned charging station), we design this migration plan that can "compensate" such migrated EVs' drivers. We define a rather general compensation function that can be customized in practice, and show how our design works under two example compensation functions, including (*i*)minimizing the detour distance for any migrated EV, and (*ii*) re-scheduling a migrated EV to a "hotspot" charging location where demanding service requests often occur in nearby locations according to historical data (e.g., downtown city area).

REC achieves this goal through two steps. In the first step, REC identifies all feasible migration plans under which the resulting modified charging scheduling decisions still guarantee predictable and tightly bounded waiting times for all EVs in the system. During this step, we will construct an auxiliary directed graph to transform the EV migration problem into a classical directed cycle detection problem. We can thus apply a rich set of existing fast algorithms designed for the directed cycle detection problem to resolve the original migration problem [6], [9], [12], [24] (e.g., using the Tarjan's strongly connected components algorithm [22], which has been shown to be very efficient in terms of runtime complexity). Then in the second step, REC further identifies the best migration plan among all feasible ones given a user-defined compensation function.

Step 1: Identify feasible migration plans.

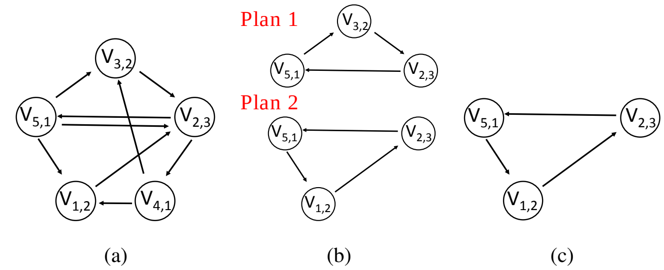

Fig. 7: An intuitive example illustrating various feasible migration plans: (a) the constructed auxiliary graph, (b) a feasible migration plan, (c) a most efficient migration plan.

In order to efficiently identify all feasible migration plans at runtime, we develop a fast graph-based method. Before presenting the detailed techniques, we show a simple example to illustrate the intuitive idea behind this method.

Example 1. *Consider an example system where five EVs are scheduled among three charging stations. When T_5's energy drops to its charging threshold, REC schedules it to s_1. Assume that T_5 requests to charge at s_2. If T_5 travels to s_2 directly, it may either experience a long waiting time or will cause later arrived taxis scheduled by REC to experience longer waiting times. Thus, an intuitive idea to resolve this is that, if T_5 travels to s_2, then one of the EVs assigned to s_2 shall be rescheduled to s_1. Such a "switching" action may make the utilizations of charging stations balanced again. If there is an EV originally scheduled to s_2 has enough energy to travel to s_1, then the problem gets easily solved. Now consider the harder case. Assuming that no EVs originally scheduled to s_2 have enough energy to travel to s_1; while T_1, which is originally scheduled to s_2, has enough energy to travel to s_3, and T_2, which is originally scheduled to s_3, has enough energy to travel to s_1. In this case, a possible migration plan for scheduling T_5 to its preferred charging station at s_2 is to re-schedule T_5 to charge at s_2 instead of s_1, T_1 to charge at s_3 instead of s_2, and T_2 to charge at s_1 instead of s_3. After these switch actions, the number of EVs scheduled to each charging station remains the same, which implies that the utilization among charging stations is still balanced.*

To find the possible migration plan for an EV in a general case, we construct an auxiliary directed graph $G(V, E)$ to transform the EV migration problem into a classical directed cycle detection problem in this constructed auxiliary graph. By solving the directed cycle detection problem using existing algorithms, we can identify all feasible migration plans that guarantee predictable waiting times. To construct this graph, for each EV $T_i \in T$, there is a corresponding vertex $v_{i,j} \in V(G)$ where $1 \leq i \leq n$ and $1 \leq j \leq m$. $v_{i,j}$ represents that EV T_i is assigned to charging station s_j. Let $t_{i,j}$ denote the time instant when T_i will arrive at s_j according to the scheduling decision given by REC. For any two different vertices v_{i_1,j_1} , $v_{i_2,j_2} \in V(G)$, there is a directed edge from v_{i_1,j_1} to v_{i_2,j_2} (e.g., $(v_{i_1,j_1}, v_{i_2,j_2}) \in E(G)$) if T_{i_1} has enough battery energy to travel to s_{j_2} by t_{i_2,j_2}.

Example 2. *We now use the same example taxi system given in Example 1 to illustrate how to construct the auxiliary graph and how to find feasible migration plans by detecting directed circles in this graph. Suppose T_1, T_2, T_3, and T_4 are dispatched to s_2, s_3, s_2, and s_1 respectively. When T_5's energy*

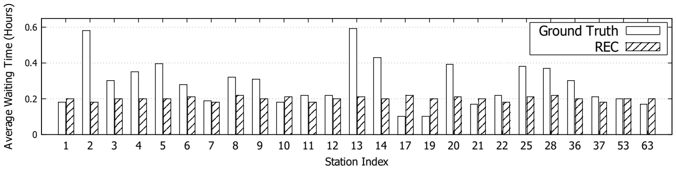

Fig. 8: Average waiting time when charging at different charging stations.

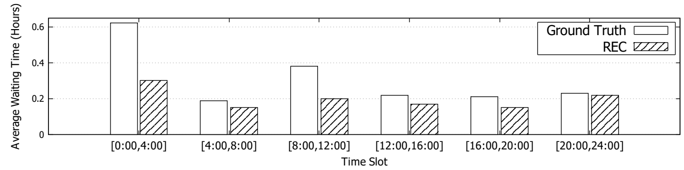

Fig. 9: Average waiting time when charging at different charging stations

drops to its charging threshold, REC schedules it to s_1. Since T_1 is scheduled to s_2, according to the definition of vertex, we have a vertex $v_{1,2}$ in the graph. For the same reason, there are 5 vertices in the graph, which are $v_{1,2}$, $v_{2,3}$, $v_{3,2}$, $v_{4,1}$ and $v_{5,1}$. Suppose T_4 is traveling to s_1 and it has enough energy to travel to s_2. According to the definition of links, there are two links in the graph, i.e., $(v_{4,1}, v_{3,2})$ and $(v_{4,1}, v_{1,2})$. Generally speaking, if there is a link from v_{i_1,j_1} to v_{i_2,j_2}, then T_{i_1} has enough energy to travel to s_{j_2}. All the vertices and links are shown in Fig. 7(a).

In the following, we illustrate how to use the graph to find a possible migration plan for T_5 if it requests to be charged at its preferred charging location s_2. Assume T_5 still has enough energy to travel to any of the four charging stations in this example system. This implies that there is a link from $V_{5,1}$ to any vertex in the graph except $V_{4,1}$, since T_4 is scheduled to the the same destination s_2 with T_5. In this graph, REC can find two directed cycles in the auxiliary graph containing the links from T_5 to s_2. They are $V_{5,1} \to V_{3,2} \to V_{2,3} \to V_{5,1}$ and $V_{5,1} \to V_{1,2} \to V_{2,3} \to V_{5,1}$. These two detected directed cycles correspond to two feasible migration plans, as shown in as illustrated in Fig. 7(b). Consider $V_{5,1} \to V_{1,2} \to V_{2,3} \to V_{5,1}$ for example. According to this directed cycle, the corresponding migration plan will be: T_5 will charge at s_2, and T_1 is re-scheduled to charge at s_3 instead of s_2, and T_2 travels to charge at s_1 instead of s_3.

Note that under this graph-based approach, for any identified feasible migration plan, the number of EVs scheduled to each charging station remains the same. Thus, if REC is able to guarantee predictable waiting times and balanced utilization among charging stations as shown in Sec. 3, then the utilization of all charging stations remains spatially and temporally balanced under any feasible migration plan. This has also been proved by our extensive experimental results as discussed in Sec. 5.

Step 2: Select the best migration plan among all feasible ones.

As discussed above, REC finds a migration plan in order to satisfy certain drivers' "greedy" requests while trying to maintaining the balancing status w.r.t. charging station utilization, through migrating some other taxis to charging locations different from the originally scheduled ones. This may have negative impact on such migrated taxis, e.g., some taxis may be already on the way to its originally scheduled charging station and has to be detour and spend additional time in order to reach the new location. Given practical considerations, REC is designed to "compensate" such migrated taxis by choosing the best migration plan that may benefit a migrated taxi the most. Intuitively, there exist a number of compensation goals that may be attractive to such migrated taxis. In the following, we use two intuitive compensation goals to illustrate how REC would select the best migration plan accordingly.

Minimum detour distance: The most compensation goal that would attract taxi drivers might be to provide a migration plan that results in the minimum detour distance, which is defined to be the distance between the newly assigned charging location and the location of the taxi when it receives the migration request. Thus, towards this goal, REC will choose the migration plan among all feasible ones that result in the minimum total detour distance considering all migrated EVs. For instance, consider the two feasible migration plans given in Example 2. If the migration plan $V_{5,1} \to V_{3,2} \to V_{2,3} \to V_{5,1}$ results a shorter total detour distance for T_2 and T_3, this migration plan will be selected. Clearly, REC can efficiently select the best migration plan in linear time complexity.

Maximum short-term revenue: Another intuitive compensation goal that may attract those migrated taxi drivers is to re-schedule them to "hotspot" charging locations where demanding service requests often occur in nearby locations according to historical data (e.g., downtown city area). Towards this goal, it is clear that REC can also easily pick the best migration plan that results in the maximum number of migrated taxis to be re-scheduled to such hotspot charging locations.

There may exist a large number of options defining the compensation function. We note that the key challenge is to identify all feasible migration plans. Whenever a specific compensation function is defined, REC can always find the best migration plan accordingly in an efficient manner.

5 Experiments

To evaluate the efficacy of REC in practice, we conducted extensive experiments based on our accessible real-world datasets as described in Sec. 2.1.

5.1 Experiment overview

To show the effectiveness of REC, we compare the performance of REC with "Ground Truth," which represents the original GPS traces analyzed from the datasets without any modifications. We also include an optimal yet impractical Oracle solution, which performs an offline iterative search that identifies the best possible charging scheduling plan that minimizes the total waiting time. Thus, Oracle clearly yields a performance upper bound, thus being useful to help us understand the performance gap between REC and an optimal solution.

These three approaches are evaluated using three practical real-world metrics that matter to any commercial public taxi fleet systems, including: (*i*) average waiting time among EVs in the fleet, (*ii*) total number of EVs that have enough energy to serve trip requests during certain time windows, and (*iii*) total travel distance of all EVs within time windows of fixed lengths.

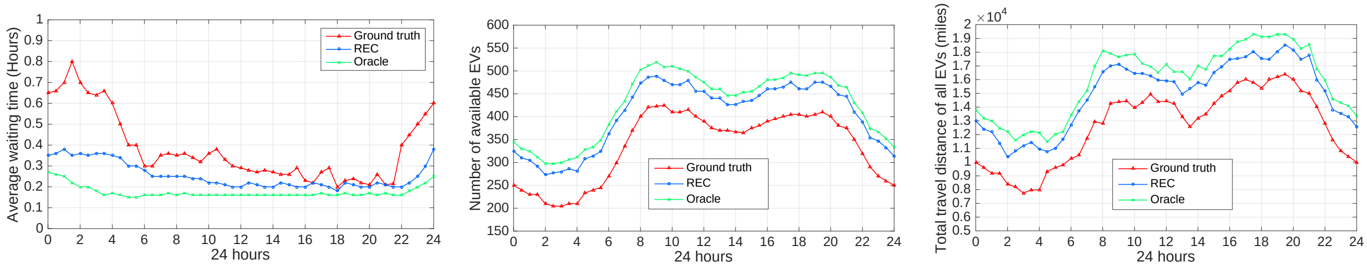

(a) Average waiting time.

(b) Total number of EVs that can serve trip requests.

(c) Total traveling distance.

Fig. 10: Serving capacity.

We have also conducted experiments to validate the second layer of REC, where minimum detour distance and maximum short-term revenue are used as the compensation goal to evaluate whether REC is able to identify good migration plans, given arbitrary drivers' requests to charge at their self-selected charging locations.

5.2 Waiting Time and Utilization Balancing

Fig. 8 shows the average waiting time at each charging station under REC and Ground Truth. As seen in this figure, the average waiting time yielded by Ground Truth varies significantly at different charging stations. For instance, the average waiting times at stations 2, 13, 14, 20 are considerably longer than other stations. On the other hand, under REC, the average waiting time is around 0.2 hours for almost all charging stations with a almost negligible variance. This implies that REC is able to provide predictable and tightly bounded waiting times as well as spatially balanced utilization among all charging stations.

Fig. 9 shows the average waiting time during six time slots (each covering four hours of a day) under REC and Ground Truth. Again, temporally speaking, Ground truth yields an average waiting time that varies significantly during different time slots. While under REC, the variation of the average waiting time is rather small. We can see that during most of the time slots, the average waiting time is around 0.2 hour except for the time slot $[0:00, 4:00)$, in which an around 0.3 hour average waiting time is observed. This again implies that in addition to spatially balanced utilization goal, REC can also achieve temporally balanced utilization among all charging stations.

5.3 Serving Capacity

Fig. 10a plots the average waiting time during 24 hours of a day. As seen in the figure, both REC and Oracle outperform Ground Truth by an average waiting time decrement of 38% and 43%, respectively, confirming the efficacy of our designed REC. Also, the performance gap between REC and Oracle is reasonably small, around 5% on average. Thus, REC may achieve a near-optimal charging scheduling plan with rather low runtime complexity.

Fig 10b plots the number of EVs that have enough energy to serve trip requests during 24 hours of a day. As seen in this figure, REC outperform Ground Truth by 32% on average. This implies that through carefully scheduling charging requests among charging stations in a temporally and spatially balanced manner, REC is able to allow more EVs to have enough energy to serve trip requests, instead of spending a long time waiting at charging stations. Moreover, compared to Oracle, REC is only being outperformed by around 5% on average. An interesting observation is that the number of available EVs increases during time window [6:00, 20:00] under all three approaches. This is because all three approaches choose to schedule more EVs to be charged during non-rush hours, i.e., [0:00, 6:00] and [22:00, 24:00].

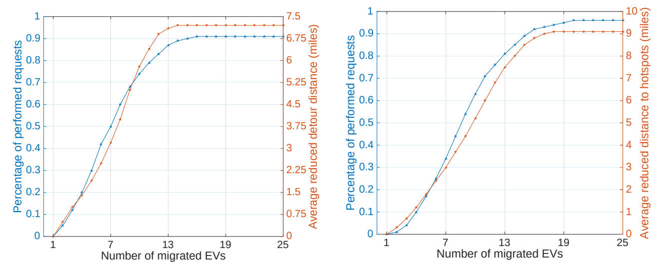

(a)

(b)

Fig. 11: Efficacy of REC when handling the practical issue using two compensation functions: (a) minimum detour distance, (b) maximum short-term revenue.

We have also conducted a set of experiment evaluating the performance under the three approaches in terms of the total travel distance of all EVs during 24 hours of a day. This metric can also be used to reflect the effectiveness of each approach since a shorter waiting time experienced by a taxi often results in a longer travel distance on the road within a window of fixed length. Fig. 10c shows the result using this metric. As seen, REC outperforms Ground Truth by a considerable margin, while remaining a rather small performance gap with the optimal Oracle solution.

5.4 Migration Policy

In this set of experiments, we evaluate the efficacy of REC when handling the practical issue discussed in Sec. 4, with those two mentioned compensation functions. Fig. 11a shows the evaluation results. We randomly generate a set of charging location requests by different taxis, which are different from the originally scheduled charging locations under REC. We use the metric "average reduced detour distance" (in miles) to reflect the compensation function of minimizing total detour distance. Similarly, we use "average reduced distance to hotspots" to represent the total reduced distance driven by a migrated EV to reach a hotspot charging location compared to the case where the same EV needs to drive from its originally scheduled charging location to this same hotspot charging location, to reflect the second compensation function of maximizing short-term revenue.

Fig. 11a(a) shows the results w.r.t. the first compensation goal. As seen in the figure, with increasing number of migrated EVs, the percentages of the performed requests increase from 0% to 90%, and the the average reduced detour distance increases from 0 to 6.75 miles. This implies that the more EVs involved in the migration plan, the larger chance of performing the drivers' charging location requests under REC. More importantly, REC is able to find such a migration plan that maximize the compensation function. One interesting observation is that when the number of migrated EVs exceeds 15, both the percentage of performed requests and the average reduced detour distance stop increasing. This implies that the best migration plan w.r.t. this compensation function can always be found with 15 migrated EVs. Fig. 11b shows similar results under the compensation function of maximizing short-term revenue. Same observations are obtained in this figure. Thus, we believe that REC can effectively handle the practical isse very well, where drivers' may make runtime requests to charge at their preferred charging stations.

6 Related work

A recent set of studies have been conducted on developing smarter EV charging schedulers. They mostly target at saving charging cost by avoiding charging during the peak hours at day time but rather charging during night hours when the power price is cheap [5], [10], [14]. Smart energy control strategies are proposed for EV charging [7], [15], [20], which are based on quadratic programming aiming at minimizing the peak load and flatten the overall load profile. Taking smart grid constraints into account, novel methods of planning EV charging are given in [11], [21], which reduce the overloading in the electricity grid and improve the charging efficiency of EVs. Different from these works, this paper address a practical problem motivated by analyzing the real-world EV taxi datasets, which is to develop a EV taxi charging scheduling framework that yields analytically predictable and tightly bounded waiting times for EV taxis, and thus temporally and spatially balanced utilization among charging stations.

7 Conclusion

Motivated by observing the unpredictably long waiting times and unbalanced charging station utilization through analyzing real-world EV taxi datasets, this paper develops a smart EV taxi charging scheduling framework, REC, which analytically guarantees predictable and tightly bounded waiting times for EV taxis in a fleet system. REC is also designed to be capable of addressing practical issues where taxi drivers may request to charge at their preferred charging stations. Extensive experiments using the real-world EV taxi datasets prove the effectiveness of REC under various scenarios.

References

[1] Google geocoding api. https://google-developers.appspot.com/maps/documentation/utils/geocoder/.

[2] Open streetmap. https://www.openstreetmap.org.

[3] A simple guide to electric vehicle charging. http://www.fleetcarma.com/electric-vehicle-charging-guide/.

[4] Zheng Dong and Cong Liu. Closing the loop for the selective conversion approach: A utilization-based test for hard real-time suspending task systems. In *Real-Time Systems Symposium (RTSS), 2016 IEEE*, pages 339–350. IEEE, 2016.

[5] Melike Erol-Kantarci and T Mouftah Hussein. Prediction-based charging of phevs from the smart grid with dynamic pricing. In *Local Computer Networks (LCN), 2010 IEEE 35th Conference on*, pages 1032–1039. IEEE, 2010.

[6] Robert W Floyd. Nondeterministic algorithms. *Journal of the ACM (JACM)*, 14(4):636–644, 1967.

[7] Rui Freire, Joaquim Delgado, João M Santos, and Anibal T De Almeida. Integration of renewable energy generation with ev charging strategies to optimize grid load balancing. In *Intelligent Transportation Systems (ITSC), 2010 13th International IEEE Conference on*, pages 392–396. IEEE, 2010.

[8] Kenji Funaoka, Shinpei Kato, and Nobuyuki Yamasaki. Work-conserving optimal real-time scheduling on multiprocessors. In *Real-Time Systems, 2008. ECRTS'08. Euromicro Conference on*, pages 13–22. IEEE, 2008.

[9] Harold N Gabow and Shuxin Nie. Finding a long directed cycle. *ACM Transactions on Algorithms (TALG)*, 4(1):7, 2008.

[10] Lingwen Gan, Ufuk Topcu, and Steven H Low. Optimal decentralized protocol for electric vehicle charging. *IEEE Transactions on Power Systems*, 28(2):940–951, 2013.

[11] Hye-Jin Kim, Junghoon Lee, Gyung-Leen Park, Min-Jae Kang, and Mikyung Kang. An efficient scheduling scheme on charging stations for smart transportation. *Security-Enriched Urban Computing and Smart Grid*, pages 274–278, 2010.

[12] D Knuth. The art of computer programming 2: Seminumerical algorithms. *MA: Addison-Wesley*, page 30, 1968.

[13] Jane W. S. W. Liu. *Real-Time Systems*. Prentice Hall PTR, Upper Saddle River, NJ, USA, 1st edition, 2000.

[14] Zhongjing Ma, Duncan Callaway, and Ian Hiskens. Decentralized charging control for large populations of plug-in electric vehicles. In *Decision and Control (CDC), 2010 49th IEEE Conference on*, pages 206–212. IEEE, 2010.

[15] Kevin Mets, Tom Verschueren, Wouter Haerick, Chris Develder, and Filip De Turck. Optimizing smart energy control strategies for plug-in hybrid electric vehicle charging. In *NOMS Wksps2010*.

[16] Fei Miao, Shuo Han, Shan Lin, John A Stankovic, Desheng Zhang, Sirajum Munir, Hua Huang, Tian He, and George J Pappas. Taxi dispatch with real-time sensing data in metropolitan areas: A receding horizon control approach. *TOASE'16*.

[17] National highway Traffic Safety Administration. Reduce greenhouse gass emissions and improve vehicle fuel efficiency. http://www.nhtsa.gov/staticfiles/rulemaking/pdf/cafe/CAFE_2014-18_Trucks_FactSheet-v1.pdf.

[18] Gregory J Offer, Vladimir Yufit, David A Howey, Billy Wu, and Nigel P Brandon. Module design and fault diagnosis in electric vehicle batteries. *Journal of Power Sources*, 206:383–392, 2012.

[19] Ahmed Yousuf Saber and Ganesh Kumar Venayagamoorthy. One million plug-in electric vehicles on the road by 2015. In *Intelligent Transportation Systems, 2009. ITSC'09. 12th International IEEE Conference on*, pages 1–7. IEEE, 2009.

[20] Somayeh Sojoudi and Steven H Low. Optimal charging of plug-in hybrid electric vehicles in smart grids. In *Power and Energy Society General Meeting, 2011 IEEE*, pages 1–6. IEEE, 2011.

[21] Olle Sundström and Carl Binding. Planning electric-drive vehicle charging under constrained grid conditions. In *POWERCON'10*.

[22] Robert Tarjan. Depth-first search and linear graph algorithms. *SIAM journal on computing*, 1(2):146–160, 1972.

[23] Wikipedia. Tesla supercharger. *https://en.wikipedia.org/wiki/Tesla_station*, 2017.

[24] Raphael Yuster and Uri Zwick. Detecting short directed cycles using rectangular matrix multiplication and dynamic programming. In *SODA2004*.