



Stochastics and Statistics

Optimal routing for electric vehicle service systems

Ying-Chao Hung^{a,*}, George Michailidis^b^a Department of Statistics, National Chengchi University, Taipei 11605, Taiwan^b Department of Statistics and Computer and Information Sciences, University of Florida, Gainesville, FL 32611-8545, USA

ARTICLE INFO

Article history:

Received 10 December 2014

Accepted 8 June 2015

Available online 18 June 2015

Keywords:

Routing

Electric vehicle system

Maximum throughput

Perturbed Lyapunov function method

Sojourn time

ABSTRACT

There is increased interest in deploying charging station infrastructure for electric vehicles, due to the increasing adoption of such vehicles to reduce emissions. However, there are a number of key challenges for providing high quality of service to such vehicles, stemming from technological reasons. One of them is due to the relative slow charging times and the other is due to the relative limited battery range. Hence, developing efficient routing strategies of electric vehicles requesting charging to stations that have available charging resources is an important component of the infrastructure. In this work, we propose a queueing modeling framework for the problem at hand and develop such routing strategies that optimise a performance metric related to vehicles' sojourn time in the system. By incorporating appropriate weights into the well-known dynamic routing discipline "Join-the-Shortest-Queue", we show that the proposed routing strategies not only do they maximise the queueing system's throughput, but also significantly mitigate the vehicle's sojourn time. The strategies are also adaptive in nature and responsive to changes in the speed of charging at the stations, the distribution of the vehicles' point of origin when requesting service, the traffic congestion level and the vehicle speed; all the above are novel aspects and compatible with the requirements of a modern electric vehicle charging infrastructure.

© 2015 Elsevier B.V. and Association of European Operational Research Societies (EURO) within the International Federation of Operational Research Societies (IFORS). All rights reserved.

1. Introduction

Over the last few years a strong push is occurring to reduce the use of hydrocarbons in the transportation sector. This trend is supported by the latest advances in battery and power electronics technology, along with government mandates on energy independence and resilience, as well as an increased emphasis on a smarter infrastructure. It is strongly enabled by the introduction of electric vehicles (EVs) and their close relatives Plug-in Hybrid Electric Vehicles (PHEVs) by major car manufacturers that have drastically increased consumer choices.

According to a recent report of the [International Energy Agency \(2012a\)](#), the transportation 'accounted for 6.7 gigatons of emitted CO₂ or 22 percent of the world's total. Further, global fuel demand for transportation is projected to grow approximately 40 percent by 2035 [International Energy Agency \(2012b\)](#), driven by the rapid adoption of automobiles in the fast growing economies of China, India and more recently in the African continent. EV/PHEVs represent an innovative technology that could help address both environmental concerns and longer term reduce dependence on fossil fuels. However, fast EV adoption relies on a number of socio-economic, as well

as technological factors. Key socio-economic factors include stringent emissions regulations, rising fuel prices and financial incentive programs [OECD/International Energy Agency \(2013\)](#), while the most pressing technological one is the large scale *deployment* of an efficient and well managed charging station infrastructure. At present, there are diverging forecasts on the growth rate of the EV population [International Energy Agency \(2011\)](#), although there is consensus that it is going to represent a sizable portion (at least 7 percent in the US) of the national fleet by 2025–30. Obviously, penetration rates could be significantly higher than these estimates depending on how the aforementioned socio-economic and technological factors evolve. Further, the impact of fast adoption of EVs on utilities and on the stability of the grid is mostly dependent on regional or even local penetration rates; e.g., Southern California Edison has projected a penetration rate larger than 7 percent in its service territory already by 2020.

The key concern regarding rapid adoption of EVs by utilities is that they could have a disruptive impact on the power grid, since under Level 1 charging conditions, an EV represents a load equivalent to 50 percent of that of a house, while under Level 2 conditions it represents a 2.5-fold equivalent load. Obviously, the extent of their impact will depend on the degree and local/regional density of their penetration, charging requirements and the time of the day they are charged. However, studies have shown that a fleet of EVs can be effectively

* Corresponding author. Tel.: +88 6229387115.

E-mail addresses: hungy@nccu.edu.tw (Y.-C. Hung), gmic@mail.ufl.edu (G. Michailidis).

powered by the underutilised electric power grid during the off-peak hours with little need to increase the energy delivery capacity of the existing grid infrastructure (Taylor, Maitra, Alexander, Brooks, & Duvall, 2010) if scheduled carefully. Hence, the literature has focused on coming up with efficient schedules of charging EVs overnight, e.g. an incentive based energy consumption controlling scheme was introduced in Caron and Kesidis (2010), and a direct load control (DLC) scheme for residential energy control was discussed in Ruiz, Cobelo, and Oyarzabal (2009); Wu, Wang, and Goel (2010). To achieve a sustainable electrification of the transportation sector, a robust charging station infrastructure needs to be in place that would not interfere with regular grid operations and at the same time address EV driver's range anxiety resulting from the limited ability to recharge EVs in a time commensurable with filling the tank of a gas-powered vehicle. Specifically, UCLA Smart Grid Energy Research Center argues that if 25 percent of all vehicles were pure EVs, the current US power grid would be challenged in meeting the demand for power. For some utilities, even adding Level-2 charging infrastructure may overload distribution transformers during peak hours. On the driver's quality-of-service issue, note that charging EVs is an inherently slow process (it takes between 20 minutes to several hours to fully charge them, depending on the technology used (Lukic, Cao, Bansal, Rodriguez, & Emadi, 2008)), thus requiring careful planning and control to accommodate customers.

In this study, we focus on the development of efficient routing strategies for a general model of EV charging systems with fully random environments, so as to provide quality service to EV drivers and release their range anxiety. This directly leads to the goal of minimising the strategy's performance metrics related to "sojourn time" (including EV's travel time, waiting time for charging, and charging time), which is different from the well-known Vehicle Routing Problem (VRP) (Dantzig & Ramser, 1959; Pillac, Gendreau, Guéret, & Medaglia, 2013) and Dynamic Traveling Repairman Problem (DTRP) (Bertsimas & Ryzin, 1990, 1993) that involve design of vehicle-delivery routes from a depot to a set of scattered demand locations, so as to minimise the overall routing cost (e.g., travel time/distance) under various side constraints (Baldacci, Toth, & Vigo, 2007; Ferrucci, Block, & Gendreau, 2013; Jaillet & Wagner, 2008; Kergosien, Lenté, Piton, & Billaut, 2011; Laporte, 2007, 2009; Li, Mirchandani, & Borenstein, 2009; Pavone, Frazzoli, & Bullo, 2007; Pillac et al., 2013; Solomon, 1987). We assume that service demands are placed by the EV drivers according to a general process, and the distribution of the locations of where EVs originate is arbitrary, although its support is determined by a compact Euclidean service region. There are K charging stations located in the service region, and each of them has a known charging speed (usually in kiloWatt-hours). Upon the request of service, each EV is directed to one of the stations according to some routing policy and heads to the station at a constant speed. After entering the station, EVs queue up to be serviced and their charging times are assumed to be randomly distributed from an arbitrary distribution. We also assume that real-time communication between vehicles and the decision maker is possible (e.g., advanced mobile phones or global positioning systems (GPS) can be integrated into the vehicle devices). These flexible assumptions are suitable for real-life settings, but also require more complicated routing strategies (e.g., dynamic/online routing) for solving the designated optimisation problem.

Note that the way we formulate the problem allows us to depict the EV service system as an acyclic network with two layers of feed-forward parallel queues (see Section 2 for details), thus facilitating the development of routing strategies and fundamental analysis regarding throughput and stability. For example, it was shown that the dynamic routing strategy "Join-the-Shortest-Queue" (JSQ) along with a scheduling policy based on "maximal matching" maximises the throughput of a general acyclic network and achieves system stability defined by the uniform mean recurrence time (Hung & Michailidis,

2012). Although the EV system described here has different characteristics (see Section 2 for details), the JSQ strategy appears to retain the property of throughput maximisation (see Section 4 for details). However, to minimise a key performance metric –the vehicle's sojourn time in the system– the routing strategy must in addition be responsive to other system states and control parameters. In this work, we show that the vehicle's sojourn time can be significantly mitigated by incorporating the following two factors as "weights" into the JSQ policy: (i) the rate at which vehicles can be charged at the stations and (ii) the distance of the EV to each station upon its service request. Further, its performance is adaptive to the change of several control parameters, such as the speed of charging stations, the distribution of demand locations, the traffic congestion level, and the vehicle speed. It should be noted that the idea to integrate the elements of queue length, service rate and distance to the charging facility in the development of routing strategies for queueing systems has not been proposed in the literature, to the best of our knowledge.

The remainder of this paper is organised as follows. Section 2 depicts a general model for the EV service system with fully random environments, characterises the system's maximum throughput based on the capacity of each service station, and defines system stability via the uniform mean recurrence time property. Section 3 introduces some routing policies, some of which have been used in different settings in the literature, while others have to property of maximising the system's throughput. The performance of some of these policies is assessed numerically in Section 5. Section 4 establishes that the proposed weighted versions of JSQ policies indeed maximise the system's throughput under fairly weak stochastic assumptions on the arrival and service (charging) processes. Since appropriate dependence structures are allowed for the arrival, service, and routing processes, a technique called the "perturbed Lyapunov function method" (Hung & Michailidis, 2012; Kushner, 1967) was employed to obtain the result. Section 5 evaluates the proposed throughput-maximising policies in terms of their sojourn-time metrics under various control parameter settings through a simulation study. Some concluding remarks are presented in Section 6.

2. Model description and system stability

Suppose there are K EV charging stations placed in a compact Euclidean region $\mathcal{R} \subset \mathbb{R}^2$, whose locations are denoted by S_1, S_2, \dots, S_K . Service demands for charging are placed by EV drivers at random locations in region \mathcal{R} according to a general random process with rate λ , and we assume the demand locations follow an arbitrary distribution F over the region \mathcal{R} . Upon request of service, each EV is guided to one of the charging stations according to some routing policy π and heads to the station, for simplicity, at a constant speed $v > 0$. Note that the latter assumption is made to simplify the exposition, although in practice v can be a non-decreasing function of λ ; e.g. one may imagine that vehicles slow down when the system is heavily loaded. We also assume that assigned EV routes do not change, while the EVs are in route to the assigned station. Finally, to simplify the formulation of the problem, we assume that each charging station comprises of an infinite buffer first-in-first-out (FIFO) queue and a charger with random service times. When the EV is charged (service completed), it leaves the system immediately.

Denote the set of EV service demand times by $\mathcal{A} = \{a_1, a_2, \dots\}$ and the set of times that EVs have finished charging at the station by $\mathcal{D} = \{d_1, d_2, \dots\}$. The collection of all event times is then denoted by $\mathcal{A} \cup \mathcal{D} = \{t_1, t_2, \dots\}$, where t_i represents either the demand arrival or system departure time. If $t_i \in \mathcal{A}$, then its associated demand location is denoted by l_i , $l_i \in \mathcal{R}$. For any routing policy π , the associated routing process is given by

$$I_k(t_i) = \begin{cases} 1 & \text{if } t_i \in \mathcal{A} \text{ and the vehicle is directed to station } k, \\ 0 & \text{otherwise,} \end{cases} \quad (1)$$




Fig. 1. An illustration of the locations of three charging stations in a rectangular region \mathcal{R} and their corresponding service regions.

where $k = 1, \dots, K$. Denote the probability that an EV has the service-demand location l and is directed to station k by $P_k(l)$, $\sum_{k=1}^K P_k(l) = 1$. The routing policy π is then associated with K service regions R_1, \dots, R_K , where

$$R_k = \{l \in \mathcal{R} : P_k(l) > 0\} \tag{2}$$

represents a set of possible locations that the EVs can be directed to station k , and $\bigcup_{k=1}^K R_k = \mathcal{R}$. Fig. 1 illustrates the locations of three charging stations in a rectangular region \mathcal{R} and their corresponding service regions, where $R_i \cap R_j = \emptyset$ for $i \neq j$.

Note that given a routing policy π , if there exists at least one location l such that $P_k(l) > 0$ for each charging station k , the demand for charging stemming from arriving EVs to the system is then divided into K parallel flows passing through, respectively, the K charging stations. At any event time t_n^+ , we denote the total number of EVs in flow k by

$$Q_k(t_n) = M_k(t_n) + N_k(t_n), \tag{3}$$

where $M_k(t_n)$ represents the number of vehicles heading to, but not yet having reached the station k , and $N_k(t_n)$ represents the number of EVs waiting in station k (including the one being served). Based on the above set of notation, the process $Q_k(t_n)$ evolves according to

$$Q_k(t_{n+1}) = Q_k(t_n) + I_k(t_{n+1}) - D_k(t_{n+1}), \tag{4}$$

where $D_k(t_{n+1})$ is the number of vehicles that complete service at station k during the time interval $(t_n, t_{n+1}]$. By definition, it is clear that $D_k(t_{n+1}) = 1$ if t_{n+1} is the departure time of a vehicle from station k and $D_k(t_{n+1}) = 0$ otherwise. Let us denote the queue length vectors at time t_n^+ by $\mathbf{Q}(t_n) = (Q_1(t_n), \dots, Q_K(t_n))$, $\mathbf{M}(t_n) = (M_1(t_n), \dots, M_K(t_n))$, and $\mathbf{N}(t_n) = (N_1(t_n), \dots, N_K(t_n))$.

We now characterise the maximum possible input rate of the service system, so as to keep it stable. First, the input rate of demand for electric charging is defined as

$$\lambda = \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{i \in \mathbb{Z}_+} \mathbf{1}_{\{t_i \leq t \text{ and } t_i \in A\}}, \tag{5}$$

which is based on the property of ergodicity for long-term observation. Second, note that each divided flow k comprises of two queues in tandem, where the lengths of the first-layer (which is actually a virtual queue) and second-layer queues at time t_n are $M_k(t_n)$ and $N_k(t_n)$, respectively. Fig. 2 gives a graphical presentation of the system described above. Since there is “no delay” on the service of the vehicles in the first-layer queues (i.e., vehicles head to the designated service stations immediately after the service request), each one of the first-layer queues can be treated as a $G/G/\infty$ (but not necessarily FIFO), with the mean service time (or mean travel time) being defined as

$$\nabla_k = \frac{1}{|R_k|} \int_{l \in R_k} \frac{\|l - S_k\|}{v} P_k(l) dl, \tag{6}$$

where $|R_k|$ is the area of service region R_k and $\|l - S_k\|$ represents the distance between location l and station k . By the above definition, the




Fig. 2. A graphical presentation of the system with K flows (charging stations), each of which is comprised of two layers of queues.

mean service rate of each first-layer queue is given by

$$\Delta_k = \frac{1}{\nabla_k}, \quad k = 1, \dots, K. \tag{7}$$

Third, let $\sigma_{k,j}$ be the charging time of the j th vehicle arriving at station k . Each one of the second-layer queues is simply a $G/G/1$, with the mean charging rate being defined as

$$\mu_k = \left(\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{j=1}^n \sigma_{k,j} \right)^{-1}, \quad k = 1, \dots, K. \tag{8}$$

Note that in practice the charging time $\sigma_{k,j}$ depends on the EV battery capacity, the power consumption determined by the driver, and the output power provided by the charging station. Therefore, μ_k simply represents the charging speed of station k (e.g. a DC Quick Charger can offer more than 50 kilowatt-hours output). Finally, let

$$r_k = \lim_{n \rightarrow \infty} \frac{\sum_{i=1}^n I_k(t_i)}{\sum_{i=1}^n \sum_{k=1}^K I_k(t_i)}, \tag{9}$$

which is the long-term proportion of the input traffic flow directed to station k . It is clear that $0 \leq r_k \leq 1$ for all $k = 1, \dots, K$, and $\sum_{k=1}^K r_k = 1$. Since all input traffic can go unimpeded through the first-layer queues, the following conditions on the second-layer queues are necessary to stabilise the system:

$$r_k \lambda < \mu_k \text{ for all } k = 1, \dots, K. \tag{10}$$

Since $\sum_{k=1}^K r_k = 1$ and $0 \leq r_k \leq 1$ for all $k = 1, \dots, K$, the conditions in (10) directly introduce an upper bound on the input rate, say, $\lambda < \mu_1 + \dots + \mu_K$. The quantity $\mu_1 + \dots + \mu_K$ is known as the maximum throughput of the system in the queueing literature.

The first step here is to develop a routing policy so that system stability can be established for any input rate $\lambda < \mu_1 + \dots + \mu_K$ under fairly mild assumptions. Since no strong assumptions are made for all the involved stochastic processes, the stability is defined via the “uniform mean recurrence time” property (Hung & Michailidis, 2012; Kushner, 1967). That is, the system is said to be stable if there exists some positive integer b_0 and nonnegative real-valued function $F(\cdot)$ such that for any n and $s = \min\{t \geq t_n : |\mathbf{Q}(t_n)| \leq b_0\}$

$$E_n[s - t_n] \leq F(\mathbf{Q}(t_n)) \text{ when } |\mathbf{Q}(t_n)| \geq b_0. \tag{11}$$

Here E_n denotes the expected value conditioned on the σ -field that contains all system information up to time t_n , and $|\mathbf{Q}(t)| = \sum_{k=1}^K Q_k(t)$. This definition indicates that when $|\mathbf{Q}(t)|$ reaches a level $b_n > b_0$, the conditional expectation of the time required to return to the value b_0 (or smaller) is bounded above by a function of b_n , uniformly in n and in the past history. It should be noted that, stability defined in (11) directly implies that $P(Q_k(t) < \infty) = 1$ for all t and k (i.e. weak stability) and $\liminf_{t \rightarrow \infty} Q_k(t)/t = 0$ almost surely for all k (i.e. rate stability). A routing policy that satisfies the rate stability for

all input rates $\lambda < \mu_1 + \dots + \mu_K$ is called the *maximum throughput policy*.

3. Routing policies

We next introduce selected routing policies, some of which are popular in the literature, while others can potentially maximise the throughput of the system.

3.1. π_1 (routing to the nearest station)

The key objective of this policy is simple and straightforward – direct EVs to the nearest station upon their demand for service (Solomon, 1987). The associated routing processes are then given by

$$I_k(t_i) = \begin{cases} 1 & \text{if } t_i \in \mathcal{A} \text{ and } k = \arg \min_{j=1, \dots, K} \|l_i - S_j\|, \\ 0 & \text{otherwise,} \end{cases} \quad (12)$$

$k = 1, \dots, K$. Note that under π_1 , the service regions satisfy that $R_i \cap R_j = \emptyset$ for all $i \neq j$ and $\bigcup_{k=1}^K R_k = \mathcal{R}$ (please see Fig. 1 for an example). This also implies that for each $k \in \{1, \dots, K\}$,

$$P_k(l) = \begin{cases} 1 & \text{if } l \in R_k, \\ 0 & \text{otherwise.} \end{cases}$$

If the distribution of originating demand locations is uniform over \mathcal{R} , then we have

$$r_k = \lim_{n \rightarrow \infty} \frac{\sum_{i=1}^n I_k(t_i)}{\sum_{i=1}^n \sum_{k=1}^K I_k(t_i)} = \frac{|R_k|}{|\mathcal{R}|},$$

which are constant values. Thus, π_1 is obviously not a maximum throughput policy.

3.2. π_2 (the round robin policy)

This policy directs EVs to the stations in turn upon their demand for service (Knust & Lücking, 2009; Wu & Down, 2009). The associated routing processes are then given by

$$I_k(t_i) = \begin{cases} 1 & \text{if } t_i = a_j \text{ for some } j \leq i \text{ and } j = k \pmod{K}, \\ 0 & \text{otherwise,} \end{cases} \quad (13)$$

$k = 1, \dots, K$. Note that under π_2 , the service regions satisfy that $R_k = \mathcal{R}$ for all $k = 1, \dots, K$. Also, π_2 is not a maximum throughput policy since $r_k = \frac{1}{K}$ for all k .

3.3. π_3 (maximum service rate)

This policy directs each EV to the station having the largest service/charging rate upon its demand for service (Baras, Dorsey, & Makowski, 1985; Buyokkoc, Varaiya, & Walrand, 1985). The associated routing processes are then given by

$$I_k(t_i) = \begin{cases} 1 & \text{if } t_i \in \mathcal{A} \text{ and } k = \arg \max_{j=1, \dots, K} \mu_j, \\ 0 & \text{otherwise,} \end{cases} \quad (14)$$

$k = 1, \dots, K$. Note that under π_3 , the service regions satisfy that $R_k = \mathcal{R}$ and $R_j = \emptyset$ for all $j \neq k$. Since now $r_k = 1$ and $r_j = 0$ for all $j \neq k$, the maximum throughput is simply μ_k . Thus, π_3 is obviously not a maximum throughput policy.

3.4. π_4 (Random Routing Policies)

Let us consider the inequalities shown in (10), given the constraints that $0 \leq r_k \leq 1$ for all $k = 1, \dots, K$ and $\sum_{k=1}^K r_k = 1$. For any $\lambda < \mu_1 + \dots + \mu_K$, it is clear that there exists at least one solution (in fact, infinitely many) $(\bar{r}_1, \dots, \bar{r}_K)$ such that

$$\bar{r}_k \lambda < \mu_k \text{ for all } k = 1, \dots, K, \quad (15)$$

where $0 \leq \bar{r}_k \leq 1$ for all $k = 1, \dots, K$ and $\sum_{k=1}^K \bar{r}_k = 1$. For example, one can always choose $\bar{r}_k = (\mu_k / \sum_{k=1}^K \mu_k)$ for all $k = 1, \dots, K$. Therefore, if the input rate λ is known, each possible solution $(\bar{r}_1, \dots, \bar{r}_K)$ then constitutes a policy associated with the routing processes, say, if $t_i \in \mathcal{A}$,

$$I_k(t_i) = 1 \text{ with probability } \bar{r}_k, \quad (16)$$

$k = 1, \dots, K$; and $I_k(t_i) = 0$ otherwise. Based on the policy, each vehicle is directed randomly to a station based on the associated probabilities $\bar{r}_1, \dots, \bar{r}_K$ upon the demand of service (this is known as a *static* problem, see (Bertsimas & Simchi-Levi, 1996; Gendreau, Laporte & Séguin, 1999) for examples). Thus, if $\bar{r}_k \neq 0$ for all k , then the service regions satisfy that $R_k = \mathcal{R}$ for all $k = 1, \dots, K$. Although π_4 appears to be throughput maximising, it is not practical from the following perspectives: (i) the information of the input rate λ is required so as to solve \bar{r}_k ; (ii) the solution of $(\bar{r}_1, \dots, \bar{r}_K)$ is not unique; and (iii) it may result in bad performance with respect to delay or sojourn time metrics (see Section 5 for examples).

3.5. π_5 (join the weighted shortest queue)

This policy directs EVs to the station having the smallest weighted queue length upon their demand for service. The associated routing processes for this policy are given by

$$I_k(t_i) = \begin{cases} 1 & \text{if } t_i \in \mathcal{A} \text{ and } k = \arg \min_{j=1, \dots, K} \{w_j Q_j(t_{i-1})\}, \\ 0 & \text{otherwise,} \end{cases} \quad (17)$$

where $w_j > 0$ is an arbitrarily weight chosen by the designer, $k = 1, \dots, K$.

Note that π_5 can be viewed as the converse of the “maximal matching” discipline in the queueing literature (Armony & Bambos, 2003; Hung & Michailidis, 2012, 2008; Mekittikul & McKeown, 1996) and many of the existing policies can be included under this framework. For example, if w_j are the same for all j , then π_5 reduces to the standard JSQ policy; if $Q_j(t_i)$ are the same for all $j = 1, \dots, K$ at time t_i and $w_j = \frac{1}{\mu_j}$, then π_5 is equivalent to π_3 ; if $Q_j(t_i)$ are the same for all $j = 1, \dots, K$ at time t_i and $w_j = \|l_i - S_j\|$, then π_5 is equivalent to π_1 . However, the choice of w_j will not affect the system’s throughput and stability (see Section 4 for details).

It is also worth noting that π_4 and π_5 are both maximum throughput policies. However, due to the shortcomings of π_4 mentioned earlier, the theoretic aspects of system stability will be investigated by utilising merely π_5 in the remainder of this work.

4. System stability

We first show that policy π_5 maximises the system’s throughput and achieves the stability defined in (11). Suppose relaxed assumptions such as correlated vehicle arrival times, correlated vehicle service times, and interdependent routing processes are made. These assumptions then lead to non-Markovian dynamics for the arrival and service processes of the system, for which a more general framework provided by the perturbed Lyapunov function method (Hung & Michailidis, 2012; Kushner, 1967) will be used. Note that to utilise the perturbed Lyapunov function method, one usually starts with a standard Lyapunov function $V(\mathbf{Q}(t_n))$ and then incorporates a perturbation term $\delta V(t_n)$ into $V(\mathbf{Q}(t_n))$. By placing a reasonable bound on $\delta V(t_n)$, $V(\mathbf{Q}(t_n)) + \delta V(t_n)$ can then be used to show the desired system stability.

The following assumptions are needed to establish the stability result.

Assumption 1. There exists a nonnegative bounded function $\rho(m)$, which goes to zero as $m \rightarrow \infty$, such that

$$E_n \left| \frac{1_{\{t_{m+1} \in A\}}}{E_n[t_{m+1} - t_m]} - \lambda \right| \leq \rho(m - n) \text{ for } m \geq n. \tag{18}$$

This assumption states that given all system information up to time t_n , the conditional arrival rate converges to its long-term arrival rate as $m - n \rightarrow \infty$. It corresponds to a general condition on the arrival process, for which interdependence (whatever its nature may be) between arrival/inter-arrival times is allowed. When the future becomes less and less predictable, it is assumed that the conditional likelihood of such interdependence gets weaker.

By (18), it is clear that $E_n[t_{m+1} - t_m]$ is bounded below by a positive constant h_0 for all $m \geq n$. In addition, it is natural to assume that $E_n[a_{m+1} - a_m]$ (recall that $a_{m+1} - a_m$ is the $(m + 1)$ th inter-arrival time) is bounded above by a positive constant h_1 for all $m \geq n$. These together imply

$$h_0 \leq E_n[t_{m+1} - t_m] \leq h_1 \text{ for } m \geq n. \tag{19}$$

Assumption 2.

$$E_n |\lambda \{I_k(t_{m+1}) - r_k\}| \leq \rho(m - n) \text{ for all } k \text{ and } m \geq n. \tag{20}$$

This assumption states that given all system information up to time t_n , upon the request of service, the conditional likelihood that a vehicle will be directed to station k converges to its long-term proportion as $m - n \rightarrow \infty$. Note that this assumption also incorporates the inter-dependence structure of each routing process.

Assumption 3.

$$E_n \left| \frac{I_k(t_{m+1})}{E_n[t_{m+1} - t_m]} - r_k \lambda \right| \leq \rho(m - n) \text{ for all } k \text{ and } m \geq n. \tag{21}$$

This is basically a combination of Assumptions 1 and 2. Moreover, it allows the existence of a cross-dependence structure between the routing process and the arrival process. When the future becomes less and less predictable, it is assumed that the conditional cross-dependence gets weaker.

Assumption 4.

$$E_n \left| \frac{D_k(t_{m+1})}{E_n[t_{m+1} - t_m]} - \mu_k \right| \leq \rho(m - n) \text{ for all } k \text{ and } m \geq n. \tag{22}$$

This assumption states equivalently that given all system information up to time t_n , the mean service rate of each station over the time interval $(t_m, t_{m+1}]$ converges to its long-term rate as $m - n \rightarrow \infty$. It corresponds to a general condition on each service process, for which interdependence between service times is allowed.

We define first the individual “routing” Lyapunov function perturbation for each routing process

$$\delta V_k^r(t_n) = -2w_k Q_k(t_n) \sum_{m=n}^{n+N-1} E_n[t_{m+1} - t_m] E_n [\lambda \{I_k(t_{m+1}) - r_k\}],$$

where the value of N will be chosen in the proof shown later. We define next the “combined” Lyapunov function perturbation for each traffic flow

$$\delta V_k^c(t_n) = 2w_k Q_k(t_n) \sum_{m=n}^{n+N-1} E_n[t_{m+1} - t_m] E_n \left[\frac{I_k(t_{m+1})}{E_n[t_{m+1} - t_m]} - r_k \lambda \right].$$

We define last the “departure” Lyapunov function perturbation for each traffic flow




Fig. 3. A graphical presentation of the service region $\mathcal{R} = [0, 30] \times [0, 30]$ and the locations of three charging stations $S_1 = (5, 15 - 5\sqrt{3})$, $S_2 = (25, 15 - 5\sqrt{3})$ and $S_3 = (15, 15 + 5\sqrt{3})$.

$$\delta V_k^d(t_n) = -2w_k Q_k(t_n) \sum_{m=n}^{n+N-1} E_n[t_{m+1} - t_m] E_n \left[\frac{D_k(t_{m+1})}{E_n[t_{m+1} - t_m]} - \mu_k \right].$$

The full Lyapunov function perturbation $\delta V(t_n)$ is given by

$$\delta V(t_n) = \sum_{k=1}^K \delta V_k^r(t_n) + \sum_{k=1}^K \delta V_k^c(t_n) + \sum_{k=1}^K \delta V_k^d(t_n).$$

Suppose the standard Lyapunov function is chosen as $V(\mathbf{Q}(t_n)) = \sum_{k=1}^K w_k Q_k^2(t_n)$ and the time-dependent Lyapunov function is given by $\tilde{V}(\mathbf{Q}(t_n)) = V(\mathbf{Q}(t_n)) + \delta V(t_n)$, we then have:

Theorem 1. Under Assumptions 1-4, π_5 maximises the system’s throughput and achieves the stability defined in (11).

Proof. The proof is sketched in the Appendix. □

5. Evaluation of sojourn time metrics

In this section, we evaluate the performance of the introduced maximal throughput policies in terms of the sojourn-time metrics via computer simulation. Note that the sojourn time in this setting corresponds to the amount of time that an EV spends in the system, which by definition includes its travel time (to the station), waiting time (for charging), and charging time. The core of the simulator was written in C++, while the simulation setup is described as follows.

Let us consider a square region $\mathcal{R} = [0, 30] \times [0, 30]$ with three EV charging stations located in the vertices of an equilateral triangle, say, $S_1 = (5, 15 - 5\sqrt{3})$, $S_2 = (25, 15 - 5\sqrt{3})$ and $S_3 = (15, 15 + 5\sqrt{3})$. A graphical representation of the service system is given in Fig. 3. To reduce the number of control parameters that will influence the sojourn time performance, here we assume the inter-arrival times between service demands and the charging times at each station are independent and exponentially distributed. The service rates at three charging stations are given by $\mu_1 = 2$, $\mu_2 = 3$ and $\mu_3 = 6$ (e.g. assume there are one Slow charger, one Rapid AC charger, and one Rapid DC charger), which induce the maximal system throughput $\mu_1 + \mu_2 + \mu_3 = 11$. The system is then simulated with the Poisson arrival process of rates $\lambda = 6, 7, 8, 9, 10$ and 10.9 under various EV speeds $v = 0.1, 0.5, 2, 5$ and 10 . Except for the numerical results provided in Section 5.3.3, we also assume the service-demand locations are independent and uniformly distributed over the square \mathcal{R} , viz., F is a uniform distribution on $[0, 30] \times [0, 30]$. Note that these setups can be easily calibrated so that practical situations can be met.

For comparison purposes, the mean and 95th percentile (P95) vehicle sojourn time were recorded based on 2×10^5 collected events

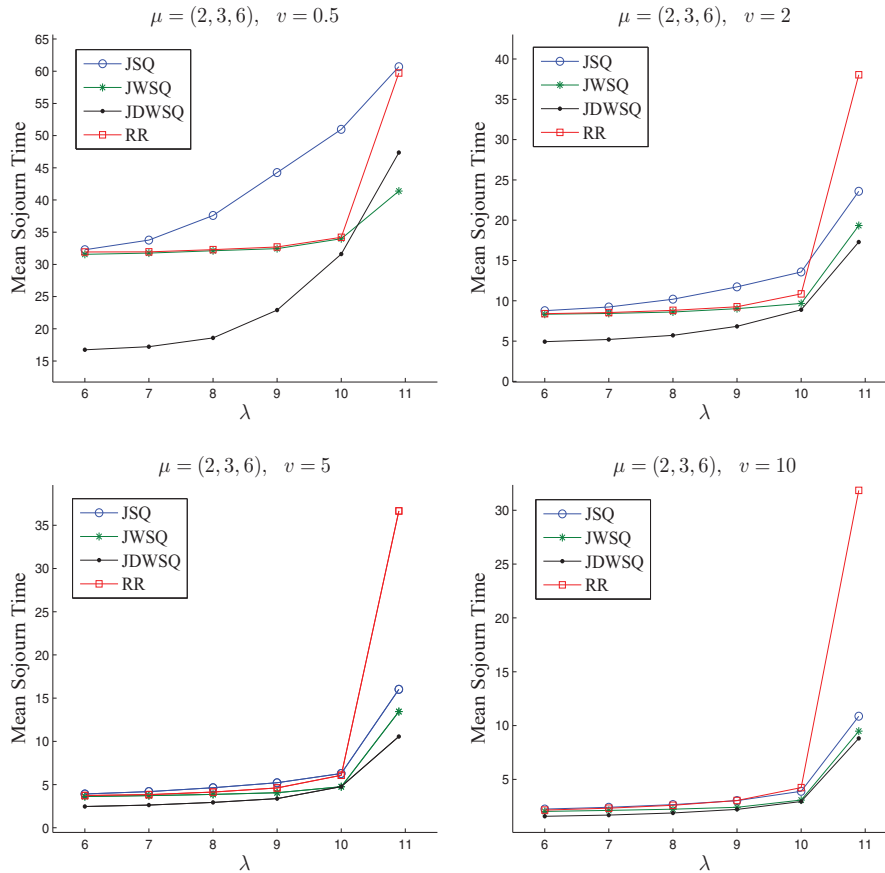


Fig. 4. The mean sojourn times under the four policies with the input rate $\lambda = 6, 7, 8, 9, 10, 10.9$ and EV speeds $v = 0.5, 2, 5, \text{ and } 10$.

(the first 10^5 events obtained from the warm-up period were excluded) under policies (i) π_4 with $(\bar{r}_1, \bar{r}_2, \bar{r}_3) = (2/11, 3/11, 6/11)$ (Random Routing policy, abbreviated as RR); (ii) π_5 with $w_j \equiv 1$ (i.e., the JSQ policy); (iii) π_5 with $w_j = \frac{1}{\mu_j}$ (Join-the-Weighted-Shortest-Queue policy, abbreviated as JWSQ); and (iv) π_5 with $w_j = \frac{\|l_i - s_j\|}{\mu_j}$ (Join-the-Doubly-Weighted-Shortest-Queue policy, abbreviated as JDWSQ). Note that w_j associated with the JDWSQ policy clearly represents a random process (not just a constant). However, the policy becomes more flexible by taking into account “the distance from the demand location to each service station”. The numerical results show that the sojourn time can be mitigated with respect to a large class of control parameters based on such a routing scheme.

5.1. The mean sojourn time

The mean sojourn times under the four policies RR, JSQ, JWSQ, and JDWSQ, for various choices of the input rate λ and EV speed v , are depicted in Fig. 4. As can be seen, the JDWSQ policy outperforms the other three policies for most of the simulation scenarios. An exception is found in the case that system is heavily loaded ($\lambda = 10.9$) and EV speed is relatively low ($v = 0.5$), for which the JWSQ policy appears to perform best in minimising the mean sojourn time. Note that as the EV speed increases, the difference between policies JWSQ and JDWSQ becomes less significant. Such a phenomenon will be further discussed in Section 5.3.1.

5.2. The 95th percentile sojourn time

The 95th percentile of the distribution of sojourn times under the four policies RR, JSQ, JWSQ, and JDWSQ, for various choices of the

input rate λ and EV speed v , are depicted in Fig. 5. As can be seen, the JDWSQ policy performs best for various choices of EV speed when the system is lightly or moderately loaded. On the other hand, the JWSQ performs best for various choices of EV speed when the system is more heavily loaded. The results may come from the intuition that, when the EV speeds are slow, the JDWSQ policy would be largely affected by the “demand locations” generated by the underlying random mechanism, thus increasing the variation of the resulting sojourn time. Analogously, the difference between policies JWSQ and JDWSQ becomes less significant as the EV speed increases.

5.3. Simulation under heavy traffic

The sojourn time performance is of particular interest in practice when the system is heavily loaded (i.e., λ is large). This usually coincides with a higher traffic congestion level and results in slower EV speeds (i.e., smaller values of v), as mentioned earlier in Section 2. In this section we introduce how to modify the JDWSQ policy so that the sojourn time metrics can be further mitigated when λ is large. We also show that the performance of the distance-based policies is fairly robust to a number of control parameters, such as the input rate λ , the EV speed v , the service rates μ_k , and the distribution F of demand locations.

5.3.1. Modification of routing policies

As can be seen from the numerical results in Section 5.1, the JDWSQ policy is not optimal in minimising the mean sojourn time when λ is large and v is small. Next, we provide insights on why policy JDWSQ may not be optimal for this metric in such a setting. Suppose there is a 2-station service system with $\mu = (6, 2)$, $Q_1(t_{n-1}) = 10$, and $Q_2(t_{n-1}) = 11$. Assume that a service demand is placed at

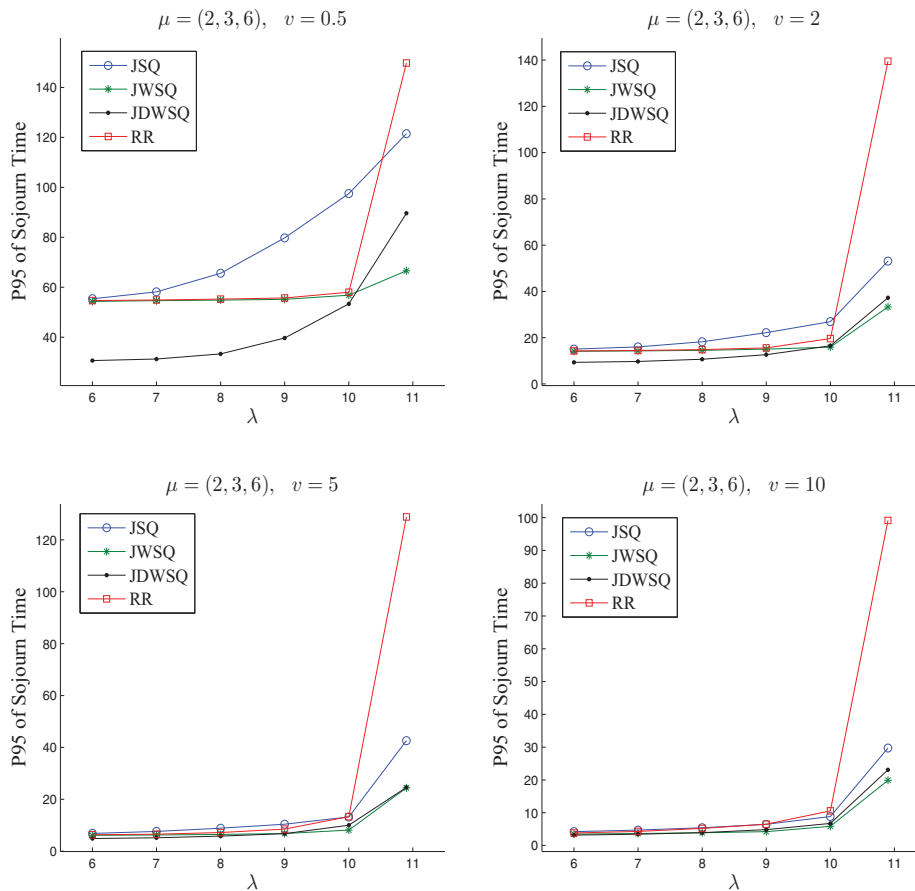


Fig. 5. The 95th percentile sojourn times under the four policies with the input rate $\lambda = 6, 7, 8, 9, 10, 10.9$ and EV speeds $v = 0.5, 2, 5, \text{ and } 10$.

time t_n with the associated location l_i such that $\|l_i - S_1\| = 30$ and $\|l_i - S_2\| = 10$. It is clear that policy JDWSQ will direct the vehicle to S_1 . Consider the scenario that the 10 vehicles on the first flow are all ahead of the current arriving vehicle (i.e., l_i is the farthest from S_1 compared to other vehicle locations on the first flow) and the 11 vehicles on the second flow are all behind the current arriving vehicle (i.e., l_i is the closest to S_2 compared to other vehicle locations on the second flow). Since v is small, it is expected that relatively few vehicles are waiting in both service stations most of the time. In this case, directing the arriving vehicle to S_2 (though it has a smaller service rate) might be a better choice since its location is much closer to S_2 and there are no vehicles ahead (thus has a shorter waiting time for service). The above example suggests that, instead of looking at all the vehicles on each flow, in the policy one may consider the number of vehicles “ahead of the arriving vehicle”. Therefore, the routing process associated with the JDWSQ policy can be justified by

$$I_k(t_i) = \begin{cases} 1 & \text{if } t_i \in \mathcal{A} \text{ and } k = \arg \min_{j=1, \dots, K} \left\{ \frac{\|l_i - S_j\|}{\mu_j} H_j(t_i) \right\}, \\ 0 & \text{otherwise,} \end{cases} \quad (23)$$

where $H_j(t_i)$ represents the number of vehicles ahead of the arriving vehicle at time t_i on the j -th flow. Note that by replacing $Q_j(t_{i-1})$ by $H_j(t_i)$ in the routing process, the JSQ and JWSQ policy can be modified in a similar way. We next evaluate the performance of these modified policies under the heavy traffic assumption.

Let us consider the same service system as in Sections 5.1 and 5.2 with $\lambda = 10.9$ (recall that the maximum throughput is 11). Tables 1 and 2 show the mean and 95th percentile sojourn times under policies JSQ, JWSQ, JDWSQ and their modifications based on $H_j(t_i)$ (denoted by JSQ*, JWSQ*, and JDWSQ*), with various choices of EV speed v . As can be seen, the modified version of the JDWSQ policy (i.e., the

Table 1

The mean sojourn times under policies JSQ, JWSQ, JDWSQ and their modifications JSQ*, JWSQ*, JDWSQ* with $\mu = (2, 3, 6)$, $v = 0.1, 0.5, 2, 5, 10$, $\lambda = 10.9$, and uniform demand locations. Note that the minimum mean sojourn time for each v is highlighted in bold.

λ	v	JSQ	JSQ*	JWSQ	JWSQ*	JDWSQ	JDWSQ*
10.9	0.1	270.64	331.20	176.32	294.64	192.35	172.07
	0.5	60.71	67.07	41.39	64.88	47.36	40.66
	2	23.59	24.84	19.34	20.81	17.31	18.13
	5	16.02	15.89	13.45	15.13	10.56	11.15
	10	10.86	10.61	9.48	12.95	8.72	8.82

Table 2

The 95th percentile sojourn times under policies JSQ, JWSQ, JDWSQ and their modifications JSQ*, JWSQ*, JDWSQ* with $\mu = (2, 3, 6)$, $v = 0.1, 0.5, 2, 5, 10$, $\lambda = 10.9$, and uniform demand locations. Note that the minimum P95 sojourn time for each v is highlighted in bold.

λ	v	JSQ	JSQ*	JWSQ	JWSQ*	JDWSQ	JDWSQ*
10.9	0.1	543	604.11	288.77	498.64	318.01	302.61
	0.5	121.48	122.29	66.62	104.68	89.60	73.02
	2	53.12	50.89	36.61	41.84	37.27	39.75
	5	42.63	44.22	24.43	33.56	24.58	26.24
	10	29.73	25.10	19.88	32.07	23.06	23.54

JDWSQ* policy) improves both the sojourn time metrics when the EV speeds are slow (i.e., $v = 0.1$ and 0.5), while policies JSQ and JWSQ do not benefit from such modifications for most of the cases. It should also be noted that, though the distance-based policies (i.e., JDWSQ and JDWSQ*) are suboptimal in minimising the 95th percentile sojourn time (see Table 2), numerical results show that they remain

Table 3

The mean and 95th percentile sojourn times under policies JSQ/JWSQ, JDWSQ, and JDWSQ* with $\mu = (4, 4, 4)$, $\nu = 0.1, 0.5, 2, 10$, $\lambda = 11.9$, and uniform demand locations.

λ	ν	JSQ/JWSQ		JDWSQ		JDWSQ*	
		Mean	P95	Mean	P95	Mean	P95
11.9	0.1	179.88	292.65	102.27	181.70	148.57	261.14
	0.5	40.17	65.61	27.30	47.39	36.31	61.72
	2	16.10	31.82	15.08	26.81	17.56	33.24
	10	7.81	16.07	6.08	13.94	9.07	22.39

Table 4

The mean sojourn times under policies JSQ, JWSQ, JDWSQ and their modifications JSQ*, JWSQ*, JDWSQ* with $\mu = (2, 3, 6)$, $\nu = 0.1, 0.5, 2, 10$, $\lambda = 10.9$, and nonuniform demand locations. Note that the minimum mean sojourn time for each ν is highlighted in bold.

λ	ν	JSQ	JSQ*	JWSQ	JWSQ*	JDWSQ	JDWSQ*
10.9	0.1	272.37	359.59	162.36	313.65	222.84	149.86
	0.5	61.35	81.89	41.81	63.27	55.59	38.91
	2	20.90	27.85	18.50	24.69	18.29	17.68
	10	13.61	21.42	13.42	15.36	11.43	8.90

fairly competitive (compared to policy JWSQ) under this particular simulation setup.

5.3.2. Homogeneous charging stations

In this section, we show that the proposed distance-based policies also work well for a system comprising of homogeneous charging stations. Consider the same service system as in Sections 5.1–5.3 with $\mu = (4, 4, 4)$, viz., the speeds of three charging stations are all the same. In this case, the two policies JSQ and JWSQ are obviously identical. Table 3 shows the mean and 95th percentile sojourn times under policies JSQ/JWSQ, JDWSQ, and JDWSQ* with a heavily loaded input rate $\lambda = 11.9$ (the maximum throughput is 12) and various EV speeds $\nu = 0.1, 0.5, 2, 10$. As can be seen from Table 3, the JDWSQ policy outperforms the JSQ/JWSQ policy in terms of both the mean and 95th percentile sojourn time for all EV speeds considered in this study. The numerical results strongly support that incorporating “distance” as another weight in the policy is beneficial for such systems. However, numerical results show that both the mean and 95th percentile sojourn time are not further improved by the JDWSQ* policy.

5.3.3. Non-uniform demand locations

In this section, we show that the distance-based policies also work well for the system with non-uniform demand locations. Let us consider the same service system in Sections 5.1–5.2 with $\lambda = 10.9$, and the distribution F of demand locations is a mixture of three uniform distributions over the rectangular regions R_1 (with vertices $(0, 0)$, $(30, 0)$, $(30, 5)$ and $(0, 5)$), R_2 (with vertices $(0, 5)$, $(30, 5)$, $(30, 10)$ and $(0, 10)$), and R_3 (with vertices $(0, 10)$, $(30, 10)$, $(30, 30)$ and $(0, 30)$). Assume that the associated weights of the three uniform distributions are given by $w_1 = 1/10$, $w_2 = 1/2$, and $w_3 = 2/5$, respectively. This then constitutes a system with half of the service demands uniformly located over R_2 (image that this is a street with extremely heavy traffic) and another half of the service demands uniformly located over R_1 and R_3 . Fig. 6 illustrates 2×10^3 demand locations that were obtained by simulating the designated mixture distribution F over the service region \mathcal{R} .

Tables 4 and 5 show the mean and 95th percentile sojourn times under policies JSQ, JWSQ, JDWSQ and their modifications JSQ*, JWSQ*, JDWSQ* with various EV speeds. As can be seen, for this particular case policy JDWSQ* outperforms all the other policies in terms of both sojourn time metrics. As the EV speed becomes slower, nu-




Fig. 6. An illustration of 2×10^3 simulated demand locations (the dots in red) based on the designated mixture distribution F over the service region \mathcal{R} . (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article).

Table 5

The 95th percentile sojourn times under policies JSQ, JWSQ, JDWSQ and their modifications JSQ*, JWSQ*, JDWSQ* with $\mu = (2, 3, 6)$, $\nu = 0.1, 0.5, 2, 10$, $\lambda = 10.9$, and nonuniform demand locations. Note that the minimum P95 sojourn time for each ν is highlighted in bold.

λ	ν	JSQ	JSQ*	JWSQ	JWSQ*	JDWSQ	JDWSQ*
10.9	0.1	557.77	642.21	264.45	505.42	468.85	251.38
	0.5	124.92	158.23	66.30	111.30	108.34	65.54
	2	55.68	57.15	43.60	56.61	53.35	41.72
	10	38.50	40.73	25.39	43.19	37.12	24.17

merical results show that the improvement in both sojourn time metrics becomes more significant.

6. Concluding remarks

We have introduced a general modeling framework for EV service systems and established that a class of dynamic routing policies maximises the throughput and maintains system stability defined via the uniform mean recurrence time under fairly weak stochastic assumptions. As the dependence structure is appropriately incorporated into the vehicle arrival, service, and routing processes, the theoretical results are established via a technique called the perturbed Lyapunov function method. Simulation results show that the proposed weighted versions of the JSQ policy, titled JWSQ, JDWSQ and JDWSQ*, outperform the existing policies in terms of both the mean and P95 of sojourn time for almost all the simulation scenarios. In addition, the performance is adaptive to the change of (i) input load λ (or traffic congestion level); (ii) EV speed ν ; (iii) service rates μ_k ; and (iv) distribution F of demand locations. It is noted that for such service systems there always exist a number of control parameters that will influence the performance measures of interest. Therefore, it would be valuable to explore the optimisation problems with respect to other control parameters, e.g. how to best allocate the charging stations so as to minimise the vehicle's mean sojourn time. Nevertheless, finding the solutions for this type of optimisation problems appears to be a challenging and involved task, and hence a topic for future work.

Acknowledgement

The authors would like to thank the Editor and two anonymous referees for their valuable comments and suggestions that help improve this manuscript. The work of George Michailidis was supported in part by NSF grants CNS 1422078 and CCF 1442493.

Appendix. Proof of Theorem 1.

For any given input rate $\lambda < \mu_1 + \dots + \mu_K$, we will show that there exists a constant $c > 0$ such that $E_n[\tilde{V}(\mathbf{Q}(t_{n+1})) - \tilde{V}(\mathbf{Q}(t_n))] \leq -c$ when $|\mathbf{Q}(t_n)|$ is large enough. This inequality together with the bound on $\delta V(t_n)$ will then imply (11). Note that

$$\begin{aligned}
 E_n[\tilde{V}(\mathbf{Q}(t_{n+1})) - \tilde{V}(\mathbf{Q}(t_n))] &= E_n[V(\mathbf{Q}(t_{n+1})) - V(\mathbf{Q}(t_n))] + \sum_{k=1}^K E_n[\delta V_k^r(t_{n+1}) - \delta V_k^r(t_n)] \\
 &\quad + \sum_{k=1}^K E_n[\delta V_k^c(t_{n+1}) - \delta V_k^c(t_n)] + \sum_{k=1}^K E_n[\delta V_k^d(t_{n+1}) - \delta V_k^d(t_n)],
 \end{aligned} \tag{24}$$

and we will examine it term by term. By mimicking the proof provided in Hung and Michailidis (2012), we can show

$$\begin{aligned}
 E_n[V(\mathbf{Q}(t_{n+1})) - V(\mathbf{Q}(t_n))] &\leq 2 \sum_{k=1}^K w_k Q_k(t_n) \{I_k(t_{n+1}) - E_n[D_k(t_{n+1})]\} + G,
 \end{aligned} \tag{25}$$

where $G = \sum_{k=1}^K w_k$;

$$\begin{aligned}
 E_n[\delta V_k^r(t_{n+1}) - \delta V_k^r(t_n)] &\leq 2w_k Q_k(t_n) E_n[t_{n+1} - t_n] \{I_k(t_{n+1})\lambda - r_k\lambda\} \\
 &\quad + 2h_1 w_k Q_k(t_n) \rho(N) + C_k^r
 \end{aligned} \tag{26}$$

for some constant $C_k^r > 0$;

$$\begin{aligned}
 E_n[\delta V_k^c(t_{n+1}) - \delta V_k^c(t_n)] &\leq -2w_k Q_k(t_n) \{I_k(t_{n+1}) - r_k\lambda E_n[t_{n+1} - t_n]\} \\
 &\quad + 2h_1 w_k Q_k(t_n) \rho(N) + C_k^c
 \end{aligned} \tag{27}$$

for some constant $C_k^c > 0$; and also

$$\begin{aligned}
 E_n[\delta V_k^d(t_{n+1}) - \delta V_k^d(t_n)] &\leq 2w_k Q_k(t_n) \{E_n[D_k(t_{n+1})] - \mu_k E_n[t_{n+1} - t_n]\} \\
 &\quad + 2h_1 w_k Q_k(t_n) \rho(N) + C_k^d
 \end{aligned} \tag{28}$$

for some constant $C_k^d > 0$.

Taking the full Lyapunov function perturbation based on (26)-(28) along with the bound given in (25), we have

$$\begin{aligned}
 E_n[\tilde{V}(\mathbf{Q}(t_{n+1})) - \tilde{V}(\mathbf{Q}(t_n))] &\leq 2E_n[t_{n+1} - t_n] \sum_{k=1}^K w_k Q_k(t_n) [I_k(t_{n+1})\lambda - \mu_k] + \varepsilon' \\
 &= 2E_n[t_{n+1} - t_n] \left(\sum_{k=1}^K \lambda w_k Q_k(t_n) I_k(t_{n+1}) - \sum_{k=1}^K \mu_k w_k Q_k(t_n) \right) + \varepsilon',
 \end{aligned} \tag{29}$$

where ε' is bounded above by

$$6h_1 \rho(N) \sum_{k=1}^K w_k Q_k(t_n) + \sum_{k=1}^K (C_k^r + C_k^c + C_k^d) + G. \tag{30}$$

Recall that for any given $\lambda < \mu_1 + \dots + \mu_K$, there always exist $\bar{r}_1, \dots, \bar{r}_K$ such that $\bar{r}_k \lambda < \mu_k$ for all $k = 1, \dots, K$, where $0 \leq \bar{r}_k \leq 1$ for

all k and $\sum_{k=1}^K \bar{r}_k = 1$. This means that there exists some constant $c_0 > 0$ such that $\bar{r}_k \lambda - \mu_k \leq -c_0$ for all $k = 1, \dots, K$. Our proof proceeds by discussing two different cases: $t_{n+1} \in \mathcal{A}$ and $t_{n+1} \notin \mathcal{A}$. Suppose now $t_{n+1} \in \mathcal{A}$ and $j = \arg \min_{k=1, \dots, K} \{w_k Q_k(t_n)\}$, then $I_j(t_{n+1}) = 1$ and $I_i(t_{n+1}) = 0$ for $i \neq j$. Thus, we have

$$\begin{aligned}
 E_n[\tilde{V}(\mathbf{Q}(t_{n+1})) - \tilde{V}(\mathbf{Q}(t_n))] &\leq 2E_n[t_{n+1} - t_n] \left(\sum_{k=1}^K \bar{r}_k \lambda w_k Q_j(t_n) - \sum_{k=1}^K \mu_k w_k Q_k(t_n) \right) + \varepsilon' \\
 &\leq 2E_n[t_{n+1} - t_n] \left(\sum_{k=1}^K \bar{r}_k \lambda w_k Q_k(t_n) - \sum_{k=1}^K \mu_k w_k Q_k(t_n) \right) + \varepsilon' \\
 &= 2E_n[t_{n+1} - t_n] \sum_{k=1}^K (\bar{r}_k \lambda - \mu_k) w_k Q_k(t_n) + \varepsilon' \\
 &\leq -2c_0 h_0 \sum_{k=1}^K w_k Q_k(t_n) + \varepsilon'.
 \end{aligned} \tag{31}$$

Note that given $|\mathbf{Q}(t_n)| > b$ for some positive constant b , we can always choose a large value of N in (30) so that

$$6h_1 \rho(N) \sum_{k=1}^K w_k Q_k(t_n) \leq c_0 h_0 \sum_{k=1}^K w_k Q_k(t_n).$$

Therefore, (31) can be expressed as

$$E_n[\tilde{V}(\mathbf{Q}(t_{n+1})) - \tilde{V}(\mathbf{Q}(t_n))] \leq -c_0 h_0 \sum_{k=1}^K w_k Q_k(t_n) + C, \tag{32}$$

where $C = \sum_{k=1}^K (C_k^r + C_k^c + C_k^d) + G > 0$. Eq. (32) directly implies

$$E_n[\tilde{V}(\mathbf{Q}(t_{n+1})) - \tilde{V}(\mathbf{Q}(t_n))] \rightarrow -\infty \text{ as } |\mathbf{Q}(t_n)| \rightarrow \infty.$$

Thus, there exist $c > 0$ and $b^* > 0$ such that

$$E_n[\tilde{V}(\mathbf{Q}(t_{n+1})) - \tilde{V}(\mathbf{Q}(t_n))] \leq -c \text{ when } |\mathbf{Q}(t_n)| > b^*. \tag{33}$$

Since we know that $|\delta V(t_n)| = O(|\mathbf{Q}(t_n)|)$, given a small $\gamma > 0$, this implies that for sufficiently large value of b^* , $|\delta V(t_n)| = |\mathbf{V}(\mathbf{Q}(t_n)) - \tilde{V}(\mathbf{Q}(t_n))| \leq \gamma \cdot \mathbf{V}(\mathbf{Q}(t_n))$. Let t_n be such that $|\mathbf{Q}(t_n)| > b^*$ and $\tau = \min\{i > n : |\mathbf{Q}(t_i)| \leq b^*\}$. By (33), we then have

$$E_n[\tilde{V}(\mathbf{Q}(t_\tau)) - \tilde{V}(\mathbf{Q}(t_n))] \leq -cE_n[\tau - n].$$

Thus,

$$\begin{aligned}
 (1 - \gamma)E_n[V(\mathbf{Q}(t_\tau))] &\leq E_n[\tilde{V}(\mathbf{Q}(t_\tau))] \\
 &\leq -cE_n[\tau - n] + V(\mathbf{Q}(t_n))(1 + \gamma) \\
 &\leq -\frac{c}{h_1} E_n[t_\tau - t_n] + V(\mathbf{Q}(t_n))(1 + \gamma),
 \end{aligned}$$

where the last inequality comes from

$$\begin{aligned}
 -cE_n[t_\tau - t_n] &= -c\{E_n[t_{n+1} - t_n] + E_n[t_{n+2} - t_{n+1}] + \dots + E_n[t_\tau - t_{\tau-1}]\} \\
 &\geq -ch_1 E_n[\tau - n].
 \end{aligned}$$

Thus,

$$E_n[t_\tau - t_n] \leq \frac{h_1}{c} \{V(\mathbf{Q}(t_n))(1 + \gamma) - 1 + \gamma E_n[V(\mathbf{Q}(t_\tau))]\}.$$

Since $V(\mathbf{Q}(t_\tau)) \leq \sup_{|\mathbf{Q}(t_i)| \leq b^*} V(\mathbf{Q}(t_i))$, this then implies the stability defined in (11). For the second case that $t_{n+1} \notin \mathcal{A}$, we have $I_k(t_{n+1}) = 0$ for all $k = 1, \dots, K$. Thus, analogous to (31), we simply obtain

$$E_n[\tilde{V}(\mathbf{Q}(t_{n+1})) - \tilde{V}(\mathbf{Q}(t_n))] \leq -2h_0 \sum_{k=1}^K \mu_k w_k Q_k(t_n) + \varepsilon'. \tag{34}$$

Since the stability defined in (11) can be shown in a similar way, the proof is then complete.

References

- Armony, M., & Bambos, N. (2003). Queueing dynamics and maximal throughput scheduling in switched processing systems. *Queueing Systems: Theory and Applications*, 44(3), 209–252.
- Baras, J. S., Dorsey, A. J., & Makowski, A. M. (1985). Two competing queues with linear costs: The μ -rule is often optimal. *Advances in Applied Probability*, 17(1), 186–209.
- Baldacci, R., Toth, P., & Vigo, D. (2007). Recent advances in vehicle routing exact algorithms. *4OR: A Quarterly Journal of Operations Research*, 5, 269–298.
- Bertsimas, D. J., & Ryzin, G. (1990). A stochastic and dynamic vehicle routing problem in the euclidean plane. *Operations Research*, 39, 601–615.
- Bertsimas, D. J., & Ryzin, G. (1993). Stochastic and dynamic vehicle routing with general demand and interarrival time distributions. *Advances in Applied Probability*, 25, 947–978.
- Bertsimas, D. J., & Simchi-Levi, D. (1996). A new generation of vehicle routing research: robust algorithms, addressing uncertainty. *Operations Research*, 44(2), 286–304.
- Buyukkoc, C., Varaiya, P., & Walrand, J. (1985). The $c\mu$ rule revisited. *Advances in Applied Probability*, 17(1), 237–238.
- Caron, S., & Kesidis, G. (2010). Incentive-based energy consumption scheduling algorithms for the smart grid. In *Proceedings of the IEEE International Conference on Smart Grid Communications* (pp. 391–396).
- Dantzig, G., & Ramser, J. (1959). The truck dispatching problem. *Management Science*, 6(1), 80–91.
- Ferrucci, F., Block, S., & Gendreau, M. (2013). A pro-active real-time control approach for dynamic vehicle routing problems dealing with the delivery of urgent goods. *European Journal of Operational Research*, 225(1), 130–141.
- Gendreau, M., Laporte, G., & Séguin, R. (1999). Stochastic vehicle routing. *European Journal of Operational Research*, 88(1), 3–12.
- Hung, Y. C., & Michailidis, G. (2012). Stability and control of acyclic stochastic processing networks with shared resources. *IEEE Transactions on Automatic Control*, 57(2), 489–494.
- Hung, Y. C., & Michailidis, G. (2008). Modeling, scheduling, and simulation of switched processing systems. *ACM Transactions on Modeling and Computer Simulation*, 18(3), 1–24. Article 12
- International Energy Agency (2012a). Keyworld energy statistics. <http://www.iea.org/publications/freepublications/publication/technology-roadmap-electric-and-plug-in-hybrid-electric-vehicles-evphev.html>.
- International Energy Agency (2012b). World energy outlook. <http://www.worldenergyoutlook.org/publications/weo-2012/>.
- International Energy Agency (2011). Technology roadmap-electric and plug-in hybrid electric vehicles.
- Jaillet, P., & Wagner, M. R. (2008). Generalized online routing: New competitive ratios, resource augmentation, and asymptotic analyses. *Operations Research*, 56, 745–757.
- Kergosien, Y., Lenté, C., Piton, D., & Billaut, J. C. (2011). A tabu search heuristic for the dynamic transportation of patients between care units. *European Journal of Operational Research*, 214(2), 442–452.
- Knust, S., & Lücking, D. (2009). Minimizing costs in round robin tournaments with place constraints. *Computers & Operations Research*, 36(11), 2937–2943.
- Kushner, H. J. (1967). *Stochastic Stability and Control*. New York: Academic Press.
- Laporte, G. (2007). What you should know about the vehicle routing problem. *Naval Research Logistics*, 54, 811–819.
- Laporte, G. (2009). Fifty years of vehicle routing. *Transportation Science*, 43, 408–416.
- Li, J. Q., Mirchandani, P. B., & Borenstein, D. (2009). Real-time vehicle rerouting problems with time windows. *European Journal of Operational Research*, 194(3), 711–727.
- Lukic, S. M., Cao, J., Bansal, R. C., Rodriguez, F., & Emadi, A. (2008). Energy storage systems for automotive applications. *IEEE Transactions on Industrial Electronics*, 55(6), 2258–2267.
- Mekkittikul, A., & McKeown, N. (1996). A starvation-free algorithm for achieving 100% throughput in an input-queued switch. In *Proceedings of the ICCCN'96* (pp. 226–231).
- OECD/International Energy Agency (2013). Global EV outlook: Understanding the electric vehicle landscape to 2020. https://www.iea.org/publications/globalviewoutlook_2013.pdf.
- Pavone, M., Frazzoli, E., & Bullo, F. (2007). Decentralized algorithms for stochastic and dynamic vehicle routing with general demand distribution. In *Proceedings of the IEEE Conference on Decision and Control* (pp. 4869–4874).
- Pillac, V., Gendreau, M., Guéret, C., & Medaglia, A. L. (2013). A review of dynamic vehicle routing problems. *European Journal of Operational Research*, 225, 1–11.
- Ruiz, N., Cobelo, I., & Oyarzabal, J. (2009). A direct load control model for virtual power plant management. *IEEE Transactions on Power Systems*, 24(2), 959–966.
- Solomon, M. M. (1987). Algorithms for the vehicle routing and scheduling problems with time window constraints. *Operations Research*, 35(2), 254–256.
- Taylor, J., Maitra, A., Alexander, M., Brooks, D., & Duvall, M. (2010). Evaluations of plug-in electric vehicle distribution system impacts. In *Proceedings of the IEEE Power and energy society general meeting* (pp. 1–6).
- UCLA Smart Grid Energy Research Center. (2010). Electric vehicle integration into the smart grid of the future - G2V & V2G. <http://smartgrid.ucla.edu>.
- Wu, R., & Down, D. G. (2009). Round robin scheduling of heterogeneous parallel servers in heavy traffic. *European Journal of Operational Research*, 195(2), 372–380.
- Wu, Q., Wang, P., & Goel, L. (2010). Direct load control (dlc) considering nodal interrupted energy assessment rate (near) in restructured power systems. *IEEE Transactions on Power Systems*, 25(3), 1449–1456.