Optimal Online Adaptive Electric Vehicle Charging

Linqi Guo
CMS, California Institute of Technology
Pasadena, CA, 91125
Email: lguo@caltech.edu

Karl F. Erliksson
Automatic Control, Lund University
Lund, Sweden
Email: tfy13ker@student.lu.se

Steven H. Low
CMS & EE, California Institute of Technology
Pasadena, CA, 91125
Email: slow@caltech.edu

***Abstract*—We propose an online linear program (OLP) based algorithm for scheduling electric vehicle (EV) charging. To determine the charging rates in each control period, OLP solves a linear program based only on EVs currently in the charging facility, assuming no future EV arrivals. We prove that OLP achieves the offline optimal where all future EV arrivals are assumed to be known in advance, provided the cost coefficients are uniformly monotone. For general cost functions, we prove that the competitive ratio is upper bounded by the variability in the cost coefficients. We demonstrate the performance of OLP using real charging data from Google and Caltech's Adaptive Charging Network.**

I. Introduction

A. Motivation and summary

We are at the cusp of a historic transformation of our energy system into a more sustainable form in the coming decades. Electrification of our transportation system will be an important component because vehicles today consume more than a quarter of energy in the US and emit more than a quarter of energy-related carbon dioxide [1]. Electrification will not only greatly reduce greenhouse gas emission, but will also have a big impact on the future grid because electric vehicles are large but flexible loads. It is widely believed that uncontrolled EV charging may stress the distribution grid and cause voltage instability, but well controlled charging can help stabilize the grid and integrate renewables, e.g., [2], [3].

We have developed and deployed an Adaptive Charging Network (ACN) at Caltech that consists of around 50 level-2 EV chargers that are capable of real-time sensing, communication, and control; see [4]. We have proposed there a simple linear program (LP) model of EV charging where, given all EVs' arrival times, departure times and energy demands over a given control horizon, our goal is to satisfy all EVs' energy demands before their deadlines without exceeding the power capacity of the electrical infrastructure. This model defines an offline LP-based algorithm which needs the full knowledge of all future EV arrivals and serves as a performance benchmark. It also motivates an online LP (OLP) algorithm that, in each control period, computes the charging rates for all existing EVs assuming there will be *no* future arrivals. An intriguing observation is made in [4] based on simulations that OLP behaves very similarly to its offline version. We review this setup in Section II.

Proc. IEEE Power and Energy Society General Meeting, Chicago, IL, USA, July 2017.
We thank Yorie Nakahira and Fengyu Zhou of Caltech for helpful discussions, and George Lee and Ted Lee of PowerFlex for providing charging data. This work has been supported by NSF grants through PFI:AIR-TT award 1602119, EPCN 1619352, CNS 1545096 CCF 1637598 and ARPA-E grant through award DE-AR0000699.

We prove in Section III that when the costs are uniformly monotone, OLP is indeed guaranteed to achieve the offline optimal even though OLP uses no knowledge about the future. We present a counterexample in Section IV that shows that OLP in general can perform arbitrarily poorly compared with offline optimal. We prove however that the competitive ratio is upper bounded by the variability of the cost coefficients. We simulate our algorithms on real charging data from Google and Caltech's ACN deployment. We present in Section V the simulation results that suggest that OLP can reduce peak demand by 30%–70%. We conclude in Section VI.

B. Related work

There is a large literature on EV charging, see [5]–[22] for example, and we can only comment on a few most related works. The benefits of controlled EV charging are well known, e.g., to reduce the need for new power plants by filling in the valleys of background demand [8], [9], [12] or help regulate frequency and renewable integration [16], [17].

Centralized algorithms are proposed in [2], [6], [12] to minimize the grid operation expenditure. Distributed algorithms are proposed in [8]–[10], [12] for demand valley filling. An online extension is proposed in [8], whose optimality, however, is not guaranteed in general. A joint optimal power flow and charging optimization problem is studied in [11]. The economic impact of EV integration on the grid are studied in [9], [12], [18].

Usually the design of adaptive charging mechanisms can be formulated as a convex optimization problem and thus be efficiently solved [8], [9], [11], [14], [15], [19]. However, these algorithms are offline in the sense that they require the information of all EVs at the beginning of scheduling. Such design is not implementable and cannot adapt to varying EV arrival patterns in real time.

Several online algorithms for different EV charging problems have been proposed prior to our work. Online scheduling of EV charging using a combination of renewable energy and energy from the grid is studied in [15]. A similar setting is formulated as a multi-processor deadline scheduling problem in [14], for which the optimal scheduler maximizes the competitive ratio against the best offline scheduler. In [19], online algorithms to minimize peak procurement from the grid incorporating predictions with uncertainty are studied.

II. Problem Setup

In this section we review the problem formulation in [4] of optimal charging and present the OLP algorithm.

Consider a facility operator that decides the optimal charging rates for a collection of EVs $\mathcal{N} = \{1, 2, \cdots, N\}$ over a finite time horizon $\mathcal{T} = \{1, 2, \cdots, T\}$. Each EV n is specified by its arrival time a_n, departure time d_n, total energy demand e_n and peak charging rate vector $\bar{\mathbf{r}}_n = (\bar{r}_{nt}, t \in \mathcal{T})$. We allow $\bar{r}_{nt}$ to depend on time t so that the information of a_n and d_n can be specified as $\bar{r}_{nt} = 0$ for $t < a_n$ or $t > d_n$. However, we assume $\bar{r}_{nt}$ is a constant over $t \in [a_n, d_n]$. Given the EV specification, a *charging profile* of EV n is a vector $\mathbf{r}_n = (r_{nt}, t \in \mathcal{T})$ of charging rates such that $0 \le r_{nt} \le \bar{r}_{nt}$ for all $t \in \mathcal{T}$ and $\sum_{t=1}^{T} r_{nt} = e_n$.

We denote by P_t the available charging power at the facility at time t. In the optimal charging problem, the operator tries to minimize a certain linear cost subject to feasibility constraints:

$$\begin{aligned}
\min_{\mathbf{r}_n : n \in \mathcal{N}} \quad & \sum_{n=1}^{N} \sum_{t=1}^{T} c_{nt} r_{nt} \\
\text{s.t.} \quad & \sum_{t=1}^{T} r_{nt} = e_n, \ n \in \mathcal{N} && \text{(1a)} \\
& \sum_{n=1}^{N} r_{nt} \le P_t, \ t \in \mathcal{T} && \text{(1b)} \\
& 0 \le r_{nt} \le \bar{r}_{nt}, \ t \in \mathcal{T}, n \in \mathcal{N} && \text{(1c)}
\end{aligned}$$

where (1a) enforces the EV charging requests being satisfied, (1b) guarantees that the total charging rate does not exceed the capacity, and (1c) enforces the rate limits[1]. This is a linear program which can be solved offline by any LP solver. This strategy of course is not implementable as it requires the knowledge of all future EV arrivals. It, however, provides an upper bound to benchmark the performance of any online algorithms. We refer to this offline LP based strategy as OPT in the sequel.

The OLP algorithm solves an online version of the optimal charging problem with all EVs that are currently in the facility, pretending there will be no future arrivals. The OLP algorithm re-solves the problem at each time slot with updated information in a way similar to model-predictive control. Specifically, for each $t \in \mathcal{T}$, denoting the remaining energy demand of EV n as e_{nt}, OLP minimizes

$$\sum_{n=1}^{N} \sum_{s=t}^{T} c_{ns} r_{ns}$$

over future charging rates $\{\mathbf{r}_{n,t+} = (r_{ns} : s \ge t)\}_{n \in \mathcal{N}}$ subject to the remaining energy demand

$$\sum_{s=t}^{T} r_{ns} = e_{nt}, \ n \in \mathcal{N}$$

together with (1b) and (1c). After the optimal rates are computed, OLP applies the optimal rates $\mathbf{r}^*_{n,t+}$ for the current time period $s = t$ to the existing EVs. OLP updates the remaining energy demands to $e_{n(t+1)} = e_{nt} - r^*_{nt}$, and repeats the cycle. [2]

Given an optimal charging problem, we denote its optimal value computed by OPT as C_{off} and the objective value achieved by OLP as C_{on}. Since any feasible solution computed by OLP is also feasible for OPT, we know in general $C_{\text{off}} \le C_{\text{on}}$. A more interesting question, which also quantifies the potential risk of using OLP, is to determine how much we lose due to the unavailability of future information. That is, we need to understand in worst case how large the competitive ratio

$$\gamma = \frac{C_{\text{on}}}{C_{\text{off}}}$$

is. It is of no surprise that the performance of OLP is closely related to the cost functions. In fact, as we will show in Section IV, if we impose no restrictions on the cost function, OLP can perform arbitrarily badly. We derive in Section III and Section IV upper bounds for the competitive ratio under assumptions that are reasonable in many practical settings. We make the following assumption throughout the paper:

Assumption. (Feasibility) The optimization problem (1) is feasible for OPT.

III. Uniformly Monotone Cost

In this section we consider a class of cost functions which encourage charging as fast as possible (see Section V-A for simulation results) and prove in this case, OLP always has the same performance as OPT whenever OLP is feasible.

Definition III.1. The cost function of an optimal charging problem is said to be *uniformly monotone* if:

1) $c_{nt} = c_{mt}$ for any $n, m \in \mathcal{N}$ and $t \in \mathcal{T}$. For such costs, we drop the subscript n and simply write $c_{nt} =: c_t$.
2) c_t is strictly increasing in t.

Our main result is summarized as Theorem III.2.

Theorem III.2. *Assume the cost function is uniformly monotone. Let $\{\mathbf{r}^{\text{off}}_n\}_{n \in \mathcal{N}}$ be an optimal solution computed by OPT, and let $\{\mathbf{r}^{\text{on}}_n\}_{n \in \mathcal{N}}$ be an optimal solution computed by OLP. Then for any $t \in \mathcal{T}$,*

$$\sum_{n=1}^{N} r^{\text{off}}_{nt} = \sum_{n=1}^{N} r^{\text{on}}_{nt} \qquad (2)$$

provided OLP is feasible at every t.

Corollary III.3. *Assume the cost function is uniformly monotone. Then we have $C_{\text{off}} = C_{\text{on}}$, i.e., OLP achieves the same performance as OPT, provided OLP is feasible at every t.*

We remark that when the cost function is uniformly monotone, there are usually infinitely many optimal solutions, for either OPT or for each step of OLP. However, Theorem III.2 tells us that the performance of OLP is independent of which exact optimal solution is computed by OLP at each time step.

[1] This model assumes the number of charging stations is sufficient to accommodate all EVs at all time, and therefore we do not consider admission control in our formulation.

[2] In principle, OLP only needs to re-solve the charging problems every time a new EV arrives. In practice, however, the rate r^*_{nt} is only a control signal to EV n that allows it to charge up to r^*_{nt} and the car may draw strictly less power for a variety of reasons. It is therefore desirable that OLP re-solves the charging problem at every t and updates the remaining energy demands using actual energy delivered $\hat{r}_{nt}$ at time t: $e_{n(t+1)} = e_{nt} - \hat{r}_{nt}$. We ignore this detail here and assume $\hat{r}_{nt} = r^*_{nt}$.

Before proving this theorem, we first characterize the optimal solutions to general optimal charging problems. The proofs of the following lemmas are provided in [23] and omitted here.

Lemma III.4. *Let* $\{\mathbf{r}_n\}_{n\in\mathcal{N}}$ *be an optimal solution to* (1) *with uniformly monotone cost. Then*

$$\sum_{n=1}^{N} r_{n1} = \min\left\{P_1, \sum_{n=1}^{N} \min\left\{e_n, \overline{r}_{n1}\right\}\right\}$$

Lemma III.5. *Let* $\{\mathbf{r}_n\}_{n\in\mathcal{N}}$ *be an optimal solution to* (1) *with uniformly monotone cost. Then for each* $t \in \mathcal{T}$*, there does not exist partial charging rates* $\{\tilde{\mathbf{r}}_{n,t-} = (\tilde{r}_{ns} : s \le t)\}_{n\in\mathcal{N}}$ *such that*

$$\sum_{n=1}^{N} \tilde{r}_{ns} = \sum_{n=1}^{N} r_{ns}, s \le t-1 \tag{3}$$

and

$$\sum_{n=1}^{N} \min\left\{e_n - \sum_{s=1}^{t-1} \tilde{r}_{ns}, \overline{r}_{nt}\right\} > \sum_{n=1}^{N} \min\left\{e_n - \sum_{s=1}^{t-1} r_{ns}, \overline{r}_{nt}\right\} \tag{4}$$

Proof of Theorem III.2. We only sketch the proof here in light of space limitation; see [23] for details.

The proof is by induction. Consider $t = 1$. Put

$$\overline{R}_1 := \min\left\{P_1, \sum_{n=1}^{N} \min\left\{e_n, \overline{r}_{n1}\right\}\right\}$$

to be the maximum aggregate charging rate. Applying Lemma III.4 to the optimal charging problem (1), we conclude $\sum_{n=1}^{N} r_{n1}^{\text{off}} = \overline{R}_1$. Applying Lemma III.4 to the optimal charging problem (1) where all EVs with $a_n \ge 2$ are removed, we conclude $\sum_{n=1}^{N} r_{n1}^{\text{on}} = \overline{R}_1$. As a result, we know (2) is true for $t = 1$.

Now assume (2) holds for any $s < t$. Define

$$\overline{R}_t^{\text{off}} := \min\left\{P_t, \sum_{n=1}^{N} \min\left\{e_n - \sum_{s=1}^{t-1} r_{ns}^{\text{off}}, \overline{r}_{nt}\right\}\right\} \tag{5}$$

and

$$\overline{R}_t^{\text{on}} := \min\left\{P_t, \sum_{n=1}^{N} \min\left\{e_n - \sum_{s=1}^{t-1} r_{ns}^{\text{on}}, \overline{r}_{nt}\right\}\right\} \tag{6}$$

to be the maximum aggregate charging rate for OPT and OLP, respectively. By applying Lemma III.4 and Lemma III.5 properly, we conclude

$$\sum_{n=1}^{N} r_{nt}^{\text{off}} = \overline{R}_t^{\text{off}} = \overline{R}_t^{\text{on}} = \sum_{n=1}^{N} r_{nt}^{\text{on}}$$

Therefore we see (2) is also true at t. The induction step is then justified. □

IV. General Cost Function

In this section we consider general cost functions and give a sequence of problem instances with $\gamma \to \infty$. This implies that OLP can perform arbitrarily badly in the full generality of cost functions. We then devise a general performance bound which provides satisfactory guarantees when the problem data exhibits enough regularity.

Example 1. In this example, we construct a sequence of problem instances where the EV characteristics are very regular while OLP can still perform arbitrarily badly. We fix an integer $L < T$ and consider the following cost function:

$$c_t = \begin{cases} 2 & t < L \\ 1 & L \le t < T \\ l & t = T \end{cases}$$

where l is a positive integer to be decided and assume there are $T - L + 1$ many EVs. EV i arrives at $t = i$ and leaves at $t = i + L - 1$ so is present for L time slots. For simplicity, we choose $e_n = P_t = \overline{r}_{nt} = 1$ for all n, t.

When EV $n = 1$ arrives, OLP would assign it to the slot $t = L$. When EV $n = 2$ arrives, OLP would put the EVs $n = 1$ and $n = 2$ to $t = L$ and $t = L+1$ in some order. Following this process, it is clear by the time EV $n = T - L + 1$ arrives, the only possible choice is to assign it to $t = T$. This all together incurs cost $C_{\text{on}} = T - L + l$. For the offline algorithm, EV $n = 1$ can simply be assigned to slot $t = 1$, so that all the remaining EVs can be shifted ahead by one time slot. This gives $C_{\text{off}} \le T - L + 2$ and we have

$$\liminf_{l\to\infty} \gamma_l = \liminf_{l\to\infty} \frac{C_{\text{on}}}{C_{\text{off}}} \ge \lim_{l\to\infty} \frac{T - L + l}{T - L + 2} = \infty$$

□

This example tells us even if the EV characteristics are simple enough, it is still possible that OLP performs poorly when there is no restriction on the cost functions. This motivates us to define the following regularity metric.

Definition IV.1. The *price variation* of a cost function is defined to be

$$V = \frac{\max_{n,t} c_{nt}}{\min_{n,t:c_{nt}>0} c_{nt}}$$

The price variation provides an upper bound on the OLP competitive ratio. The following proposition is proved in [23].

Proposition IV.2. *For any cost function, we have:*

$$\gamma \le V$$

provided OLP is feasible at every t.

This bound is tight in the sense that we can design a sequence of problem instances whose competitive ratio approaches V. It implies that when the cost coefficients do not vary too much, for instance when c_{nt} are retail electricity prices for which V is typically small, the performance of OLP is guaranteed to be close to offline optimal.

TABLE I
POWER SAVINGS FROM OLP WITH DIFFERENT COST FUNCTIONS.

Cost Function	**OLP Power Savings**		
	MTV(2016-01-14)	Caltech(2016-02-17)	Caltech(2016-03-01)
1	40%	55%	69%
2	40%	55%	69%
3	36%	42%	55%

V. EVALUATION

In this section, we present simulation results demonstrating the performance of OLP using real charging data. The majority of the data is from Google's charging facilities in Mountain View and Sunnyvale, and we also include about three months of data collected from Caltech's ACN deployment[3]. The overall data covers more than 52,000 charging sessions from 104 different locations over more than 4,000 charging days. We examine the power capacity savings by applying OPT and OLP over the data sets. That is, we would like to determine the minimal P such that OPT or OLP is still feasible when we choose $P_t = P$ for all $t \in \mathcal{T}$.[4] The power savings is defined to be the ratio between saved power capacity and original peak power without adaptive charging. We demonstrate that in the case of uniformly monotone costs, OLP achieves the same performance as OPT up to negligible numerical errors in majority cases, and the growth of such error with respect to system load is moderate. This confirms Corollary III.3 and demonstrates the numerical stability of OLP in practice.

A. *Impact of cost function*

For the OPT algorithm, the cost function does not affect the feasible region of the charging problem and thus any choice of cost function would lead to the same power savings. This, however, is not true in general for OLP due to the possibility of myopic assignment to future slots, as demonstrated by Example 1. It is thus of interest to see how the different choice of cost functions impact the performance of OLP. We mainly examine three cost functions:

1) $c_{nt} = t\,(1 - \ell_{nt})$, where ℓ_{nt} is the laxity defined as

$$\ell_{nt} = 1 - \frac{e_{nt}}{\sum_{s \geq t} \overline{r}_{ns}}$$

This metric measures how flexible the EV is in terms of future scheduling and has been widely studied for example in [15]. This cost function encourages charging as fast as possible and prioritizes EVs with low laxity.
2) $c_{nt} = t$. This cost function only encourages charging as fast as possible.
3) $c_{nt} = 1$. This cost function does not distinguish current and future slots. OLP is free to choose any feasible point at each step.

The power savings for three sets of charging data are shown for the different choices of cost functions in Table I. One can see that cost functions 1 and 2 perform equally and are superior to cost function 3. It is also observed in the simulation that cost function 1 typically creates oscillatory charging rates due to the tendency to equalize the laxity of different EVs. Such oscillatory behaviour is generally undesirable for charging stations and EVs. Consequently, cost function 2 is the best choice in terms of OLP performance as well as implementation based on these datasets. We thus deploy the cost function $c_t = t$ for all the following simulations.

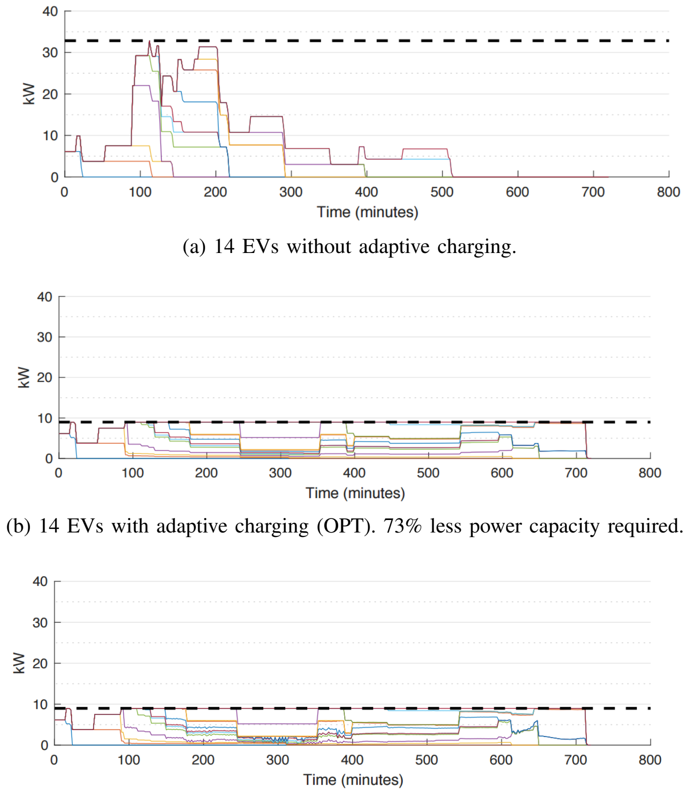

(a) 14 EVs without adaptive charging.

(b) 14 EVs with adaptive charging (OPT). 73% less power capacity required.

(c) 14 EVs with adaptive charging (OLP). 73% less power capacity required.

Fig. 1. Example profiles illustrating the power savings with and without adaptive charging algorithms.

B. *Power capacity savings*

We now demonstrate the power capacity savings through coordinated adaptive charging compared to charging EVs individually, and illustrate that OLP achieves the same performance as OPT when the cost function is uniformly monotone. In Fig. 1, the aggregate power consumption of 14 EVs as a function of time is shown for the case without adaptive charging, with OPT adaptive charging and with OLP adaptive charging, respectively. In this particular case, one can see that 73% less power capacity is required to meet the energy demand if adaptive charging is applied, regardless of whether OPT or OLP is used. Adaptive algorithms manage to schedule the charging rates so that the power consumption is more constant over time so that the peak power is significantly reduced. It can also be observed that the *aggregate* charging rates are always the same between OPT and OLP at each time slot, confirming Theorem III.2, even though the individual charge rates r_{nt}^{off} and r_{nt}^{on} may differ, e.g., between 600 mins and 700 mins in Fig. 1(b) and (c).

[3]The Google data as far is confidential. For ACN data, real-time statistics is available at http://ev.caltech.edu/. The ACN deployment is not fully implemented yet and we rely on numerical simulations here.

[4]This can be done using bisection search by making P an optimization variable.

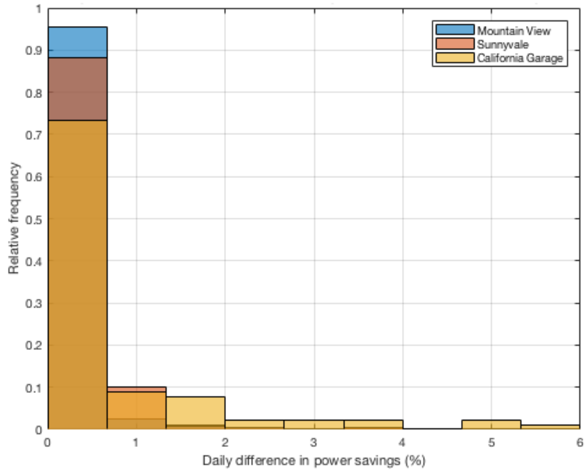

Fig. 2. Histogram of difference in daily power savings between OLP and OPT.

TABLE II
OFFLINE-ONLINE PERFORMANCE GAP UNDER DIFFERENT SYSTEM LOADS.

Number of EVs	**Average difference in power savings**		
	Caltech Campus	Mountain View	Sunnyvale
> 0	0.72%	0.078%	0.14%
> 5	0.84%	0.12%	0.23%
> 10	0.89%	0.12%	0.23%
> 15	0.96%	0.15%	0.24%
> 20	0.96%	0.18%	0.17%

C. *Offline-online performance gap*

To assess the overall performance of OLP against OPT, we examine the histogram of the difference in daily power capacity savings obtained by running OPT and OLP, as shown in Fig. 2. In this figure, all days from the data set were included and the histogram has been normalized so that the height of the bins sums to 1 for each data set. One can see from the figure that in majority of the data sets, OLP and OPT performed equally, and the performance gap is well bounded by 2% over 95% of the whole data set. This verifies the optimality of OLP as shown in Corollary III.3[5].

Intuitively, OLP is more prune to such numerical errors when the system is heavily loaded. It is therefore useful to see how different system load levels impact the performance degradation of OLP. In Table II, the average daily difference in power savings for OPT and OLP are calculated and categorized by the number of EVs that was connected to the system. The first row includes all days, while the second row only includes days with number of EVs greater than 5, etc. One can see from the simulation results that the performance degradation of OLP with respect to system load is very moderate.

VI. CONCLUSION

In this paper, we study an online EV charging algorithm based on linear programming (OLP) in terms of its performance against offline optimal for different cost functions. Real charging data from Google and Caltech Adaptive Charging Network is used to demonstrate the performance of OLP for uniformly monotone costs and illustrate the numerical stability of OLP in practical settings.

It is interesting to see whether the myopic decision made by OLP can be improved by incorporating future predictions, say through Bellman equation in an infinite horizon setting. OPT has access to exact future predictions, while the prediction for OLP can suffer from infinitely large errors. They thus form the extreme points on the spectrum of model prediction uncertainty. We are still investigating the fundamental online-offline gap or the average-case performance degradation when predictions of different level with uncertainty are available. The feasibility problem for OLP under different cost functions is also an important part for our future work.

[5]Noticeable performance gap appears when our assumption about the OLP feasibility is violated.

REFERENCES

[1] United States Department of Transportation, "National transportation statistics," https://www.rita.dot.gov/bts/sites/rita.dot.gov.bts/files/publications/national_transportation_statistics/index.html.

[2] J. Lopes, F. Soares, and P. Almeida, "Integration of electric vehicles in the electric power system," *Proceedings of the IEEE*, vol. 99, no. 1, pp. 168–183, Jan 2011.

[3] M. Caramanis and J. M. Foster, "Management of electric vehicle charging to mitigate renewable generation intermittency and distribution network congestion," in *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on*, Dec 2009, pp. 4717–4722.

[4] G. Lee, T. Lee, Z. Low, S. H. Low, and C. Ortega, "Adaptive charging network for electric vehicles," in *Proceedings of the IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, December 2016.

[5] L. Gan, U. Topcu, and S. H. Low, "Optimal decentralized protocol for electric vehicle charging," in *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*. IEEE, 2011, pp. 5798–5804.

[6] E. Sortomme, M. M. Hindi, S. D. J. MacPherson, and S. S. Venkata, "Coordinated charging of plug-in hybrid electric vehicles to minimize distribution system losses," *IEEE Transactions on Smart Grid*, vol. 2, no. 1, pp. 198–205, March 2011.

[7] C. Roe, F. Evangelos, J. Meisel, A. Meliopoulos, and T. Overbye, "Power system level impacts of PHEVs," *Proc. Hawaii International Conference on System Science*, pp. 1–10, 2009.

[8] L. Gan, A. Wierman, U. Topcu, N. Chen, and S. H. Low, "Real-time deferrable load control: Handling the uncertainties of renewable generation," in *Proceedings of the Fourth International Conference on Future Energy Systems*, ser. e-Energy '13. New York, NY, USA: ACM, 2013, pp. 113–124.

[9] Z. Ma, D. S. Callaway, and I. A. Hiskens, "Decentralized charging control of large populations of plug-in electric vehicles," *IEEE Transactions on Control Systems Technology*, vol. 21, no. 1, pp. 67–78, Jan 2013.

[10] L. Gan, T. Topcu, and S. H. Low, "Stochastic distributed protocol for electric vehicle charging with discrete charging rate," in *2012 IEEE Power and Energy Society General Meeting*, July 2012, pp. 1–8.

[11] S. Sojoudi and S. H. Low, "Optimal charging of plug-in hybrid electric vehicles in smart grids," in *Proc. IEEE PES General Meeting*. IEEE, 2011.

[12] K. Clement, E. Haesen, and J. Driesen, "Coordinated charging of multiple plug-in hybrid electric vehicles in residential distribution grids," in *Power Systems Conference and Exposition, 2009. PSCE '09. IEEE/PES*, March 2009, pp. 1–7.

[13] M. C. Caramanis and J. M. Foster, "Coupling of day ahead and real-time power markets for energy and reserves incorporating local distribution network costs and congestion," in *Communication, Control, and Computing (Allerton), 2010 48th Annual Allerton Conference on*, Sept 2010, pp. 42–49.

[14] S. Chen and L. Tong, "iEMS for large scale charging of electric vehicles: Architecture and optimal online scheduling," in *Smart Grid Communications (SmartGridComm), 2012 IEEE Third International Conference on*, Nov 2012, pp. 629–634.

[15] A. Subramanian, M. J. Garcia, D. S. Callaway, K. Poolla, and P. Varaiya, "Real-time scheduling of distributed resources," *IEEE Transactions on Smart Grid*, vol. 4, no. 4, pp. 2122–2130, Dec 2013.

[16] C. Quinn, D. Zimmerle, and T. H. Bradley, "The effect of communication architecture on the availability, reliability, and economics of plug-in hybrid electric vehicle-to-grid ancillary services," *Journal of Power Sources*, vol. 195, no. 5, pp. 1500 – 1509, 2010. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0378775309015195

[17] J. Lopes, P. Almeida, and A. da Silva, "Smart charging strategies for electric vehicles: Enhancing grid performance and maximizing the use of variable renewable energy sources," *Proceedings of Internationl Battery, Hybrid and Fuel Cell Electric Vehicle Symposium and Exhibition*, pp. 1–11, 2009.

[18] S. Shao, T. Zhang, M. Pipattanasomporn, and S. Rahman, "Impact of tou rates on distribution load shapes in a smart grid with phev penetration," in *IEEE PES T D 2010*, April 2010, pp. 1–6.

[19] S. Zhao, X. Lin, and M. Chen, "Peak-minimizing online ev charging: Price-of-uncertainty and algorithm robustification," in *2015 IEEE Conference on Computer Communications (INFOCOM)*, April 2015, pp. 2335–2343.

[20] M. Kefayati and C. Caramanis, "Efficient energy delivery management for phevs," in *Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on*, Oct 2010, pp. 525–530.

[21] E. Bitar and Y. Xu, "Deadline differentiated pricing of deferrable electric loads," *IEEE Transactions on Smart Grid*, vol. PP, no. 99, pp. 1–1, 2016.

[22] B. Wang, Y. Wang, C. Qiu, C. C. Chu, and R. Gadh, "Event-based electric vehicle scheduling considering random user behaviors," in *2015 IEEE International Conference on Smart Grid Communications (SmartGridComm)*, Nov 2015, pp. 313–318.

[23] L. Guo, K. Erliksson, and S. H. Low, "Journal version of the paper, under preparation."