Preprints of the 18th IFAC World Congress
Milano (Italy) August 28 - September 2, 2011

A Decentralized MPC Strategy for Charging Large Populations of Plug-in Electric Vehicles ⋆

Zhongjing Ma * **Ian Hiskens** ** **Duncan Callaway** ***

* *School of Automation, Beijing Institute of Technology, Beijing, China (e-mail: mazhongjing@bit.edu.cn)*
** *Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, USA (e-mail: hiskens@umich.edu)*
*** *Energy and Resources Group, University of California, Berkeley, USA, (e-mail: dcal@berkeley.edu)*

Abstract: This paper presents a framework for decentralized coordination of plug-in electric vehicle (PEV) charging patterns in scenarios where the future cannot be perfectly predicted. We begin with the mathematical formulation of the decentralized problem, in which individual PEVs minimize their own charging costs, which are a function of total system demand. We summarize results from our prior work in this area, relating specifically to convergence and uniqueness of "valley filling" charging strategies in situations where future system states are known with perfect accuracy. We then present an approach to manage forecast uncertainty by allowing decentralized agents to continually update their optimal control trajectories subject to revised forecasts of system states. We show that the resulting trajectories are strongly influenced by the accuracy of the forecast over the charging period.

Keywords: Predictive control; Electric vehicles; Load regulation; Nash games.

1. INTRODUCTION

Plug-in electric vehicles (PEVs) have the potential to replace a substantial portion of the current conventional petroleum-combustion vehicle fleet over the next few decades. A number of studies have been undertaken recently to explore the potential impacts of high penetrations of PEVs on the power grid, e.g. Denholm and Short [2006], Rahman and Shrestha [1993], Koyanagi and Uriu [1997], Koyanagi et al. [1999]. In most cases, these studies presume that the best way to charge PEVs is by filling the overnight "valley" in non-PEV electricity demand, such that the aggregate PEV demand together with non-PEV demand remains constant during the charging period. In these studies, the question of how vehicles might be encouraged or controlled to achieve this is not formally addressed.

In Ma and Callaway [2011], we study centralized optimal charging control for large populations of PEVs. The cost to charge at any given moment is determined by the total demand on the grid, which is the summation of the inelastic non-PEV base demand together with the total aggregated demand of the PEV population. This work rigorously explores the conditions under which the socially optimal (i.e. marginal generation cost minimizing) control strategy results in *valley filling*. The central finding is that, to guarantee valley filling, the product of total electricity demand with marginal cost must be convex.

In Ma et al. [2010, 2011] we examine the problem from a decentralized perspective, where each PEV makes its own charging decisions, subject to minimizing the product of its demand with real-time marginal electricity cost. In this case, the PEV charging control problems become a class of noncooperative games, where PEV agents share the electricity resources on a finite collection of charging instants. We present an iterative computational algorithm to identify a Nash equilibrium as follows: Before the start of the charging interval, all agents simultaneously update their individual best (or greedy) response with respect to the aggregate mass behavior; the process repeats with individuals optimizing against the results of the previous iteration. We show that this iterative process may not converge to valley filling without a quadratic penalty term (which is demonstrably small) for the deviation of the individual control strategy from the population average. However, under certain mild conditions, the decentralized charging control algorithm drives the system asymptotically to a unique Nash equilibrium that is nearly globally optimal. In the case of homogeneous PEV populations, this unique Nash equilibrium becomes a perfect valley-filling charging strategy. In a finite population of PEVs, the algorithm converges to an ε-Nash equilibrium, where ε tends to zero as the number of PEVs approaches infinity.

The decentralized charging strategies proposed in Ma et al. [2010, 2011] are effective only if available generation, non-PEV demand and electricity price over the charging

⋆ Research supported by the Michigan Public Service Commission through grant PSC-08-20, the National Science Foundation through EFRI-RESIN grant 0835995, and the Excellent Young Scholars Research Fund of Beijing Institute of Technology.

interval are known with perfect accuracy. Of course, in actuality these quantities are always subject to forecast error, and contingencies (such as the loss of a transmission line or forced generator outage) can significantly change operating conditions. As a consequence, the proposed charging strategy is not robust and likely to result in far-from-optimal charging costs for each PEV. Furthermore, it does not attempt to harness the inherent flexibility of PEV charge rates to mitigate the impact of contingencies.

Therefore, the main objective of this paper is to propose and explore the performance of a revised version of these game-based decentralized PEV charging control methods using concepts from model predictive control (MPC). At the core of the method is a local optimization scheme that generates a trajectory of control decisions. As with MPC, only the first instance of each optimal trajectory is used; subsequently, measurements are acquired, a new trajectory is computed, and a new control sequence is generated. This new sequence will differ from the previous one to the extent that measured quantities do not match their prediction from the previous time step, or in the case of a revised forecast. We note that the strategy we develop does not employ a "receding horizon", which typifies many MPC strategies. This is due to the fact that the PEV charging problem has an explicit stopping time (the end of the charging period), unlike tracking control problems which usually operate indefinitely into the future.

The paper is organized as follows: In Section 2 we review the decentralized charging control problems for large populations of PEVs studied in Ma et al. [2010, 2011]. In Section 3 we implement decentralized MPC methods for the underlying charging control problems established in Section 2. We then explore the performance of the MPC approach using numerical examples in Section 4. Section 5 provides conclusions and considers avenues for future research.

2. DECENTRALIZED CHARGING CONTROL OF LARGE PEV POPULATIONS

The paper considers strategies for controlling the charging of a significant penetration of PEVs with total population size N. Accordingly, the PEV population is denoted by $\mathcal{N} = \{1, 2, \cdots, N\}$. We assume all PEVs share the common charging interval $[0, T]$, i.e., they all begin charging at time $t = 0$, and conclude charging at the final time $t = T$.

We define the state of charge (SOC), normalized by the battery capacity, as $0 \leq x_{nt} \leq 1$, for the n-th PEV at time $t \in [0, T]$. Variation of the SOC over the charging interval $[0, T]$ is dependent upon the charging strategy $\mathbf{u}_n \geq 0$, and is described by the simplified model,

$$x_{n,t+1} = x_{nt} + \frac{\alpha_n}{\beta_n} u_{nt}, \quad x_{n0} \in [0, 1), \tag{1}$$

where β_n and α_n denote the battery size and the (constant) charger efficiency of the n-th PEV respectively.

A collection of PEV charging strategies, $\mathbf{u} \equiv (u_{nt}; n \in \mathcal{N}, t \in [0, T])$, is an *admissible full-charging control* of the PEV population, if $u_{nt} \geq 0$, and $x_{nT} = 1$, i.e. the battery of every PEV is fully charged at the terminal instant T. It follows from (1) that this terminal condition can be written $\sum_{t=0}^{T-1} u_{nt} = \frac{\beta_n}{\alpha_n}(1 - x_{n0})$ for all $n \in \mathcal{N}$. We therefore denote the set of admissible full-charging controls for the n-th PEV by

$$\mathcal{U}_n(x_{n0}) = \Big\{ u_{ns}; u_{ns} \geq 0, \sum_{s=0}^{T-1} u_{ns} = \frac{\beta_n}{\alpha_n}(1 - x_{n0}) \Big\}. \tag{2}$$

Subject to an admissible charging control $\mathbf{u} \in \mathcal{U}(\mathbf{x}_0)$, the cost associated with delivering the total system demand is given by,

$$\mathbb{J}(\mathbf{u}) = \sum_{t=0}^{T-1} p(r_t^N(\mathbf{u}_t)) \big(D_t^N + \sum_{n=1}^{N} u_{nt} \big), \tag{3}$$

where $p(r_t^N(\mathbf{u}_t))$ is the electricity charging price at instant t, and D_t^N is the total inelastic non-PEV demand at instant t. It is assumed D_t^N is unrelated to $\mathbf{u}$. We further assume the electricity charging price $p(r_t^N(\mathbf{u}_t))$ is determined by the ratio between the total demand and the total generation capacity, so that

$$r_t^N(\mathbf{u}_t) \triangleq \frac{1}{C^N} \big(D_t^N + \sum_{n=1}^{N} u_{nt} \big), \tag{4}$$

where C^N denotes the total generation capacity. An implicit assumption in (4) is that electricity price is a function of instantaneous demand only.

Centralized control of PEV charging calls for a strategy that minimizes the cost function (3). However, for large populations of PEVs, such control would require significant communications and computational capability. In order to circumvent those difficulties, we proposed in Ma et al. [2010, 2011] a novel decentralized algorithm where all PEV agents simultaneously update their individual responses with respect to an (electricity) charging price that is seen by all PEVs. The resulting iterative process takes place prior to the actual charging interval, with each PEV implementing its schedule later, during the actual charging period. This open-loop strategy cannot adjust as system conditions vary from the forecast. The paper addresses this issue by establishing an MPC form of the control algorithm. A review of the original game-based decentralized strategy is provided in the remainder of this section. The MPC adaptation of this strategy will be introduced in Section 3.

This paper considers the properties of systems where the number of PEVs is sufficiently large that the action of each individual PEV on the system is negligible, but the action of the aggregation of PEVs may be significant. We are therefore interested in the asymptotic properties of systems in the large N limit. To ensure that key properties of the system are preserved at that limit, we make the following asymptotic assumptions as PEV population size approaches infinity,

$$\lim_{N \to \infty} \frac{D_t^N}{N} = d_t, \qquad \lim_{N \to \infty} \frac{C^N}{N} = c. \tag{5}$$

The implication inherent in (5) is that larger power systems, with greater capacity and base demand, are required to support large numbers of PEVs. Direct substitution into (4) gives,

$$\lim_{N \to \infty} r_t^N(\mathbf{u}_t) = \frac{1}{c}(d_t + \overline{\mathbf{u}}_t) \triangleq r_t(\overline{\mathbf{u}}_t) \tag{6}$$

where

$$\overline{\mathbf{u}}_t \triangleq \lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} u_{nt}. \tag{7}$$

The function $r_t(\cdot)$ is implicitly dependent upon the non-PEV demand d_t. Consequently, the optimization process requires knowledge of the (future) demand over the entire charging interval. This information must come from a forecast. We will use the notation $d_{s|t}$ to denote the forecast value of demand at time s based on a forecast made at a prior time $t \le s$.

The decentralized charging strategy proposed in Ma et al. [2010, 2011] allows all agents to simultaneously update their individual response with respect to the total (PEV plus non-PEV) demand, such that their cost is minimized. Since all agents simultaneously update their strategy, resources that are cheap during one charging interval may become expensive during the next, and vice versa. This oscillating behavior can result in non-converge. In order to mitigate such oscillations, each PEV's cost function is modified to include a quadratic term that penalizes the deviation of its charging strategy from the average over the entire PEV population,

$$J_n(\mathbf{u}) = \sum_{t=0}^{T-1} \Big(p(r_t) u_{nt} + \delta \big(u_{nt} - \overline{\mathbf{u}}_t\big)^2 \Big) \tag{8}$$

with r_t defined in (6).

The decentralized algorithm for determining optimal charging controls can be summarized as follows:

(P1) The utility broadcasts a prediction of the non-PEV demand $d_{.|0}$ at initial instant 0 to all PEVs.

(P2) Each PEV proposes a charging strategy that is the solution of

$$\min_{\mathbf{u}_n \in \mathcal{U}_n(x_{n0})} J_n(\mathbf{u})$$

where the charging cost J_n given by (8) is defined with respect to $d_{.|0}$ and a common average (or aggregate) PEV demand $\overline{\mathbf{u}}$ broadcast by the utility.

(P3) The utility collects all the individual charging strategies proposed in (P2), and updates the average PEV demand strategy. This updated aggregate PEV demand is rebroadcast to all the PEVs.

(P4) Repeat (P2) and (P3) until the strategies proposed by the agents no longer change.

Some time later, when the actual charging start time occurs, each PEV implements the strategy it obtained from the (P1)-(P4) negotiations.

Under the assumption that the number of PEVs N is sufficiently large, at convergence of the iterative procedure (P1)-(P4), assuming it occurs, the resulting collection of individual charging strategies is an Nash equilibrium.

Consider the following assumptions:

(A1) The price function $p(r)$ is continuously differentiable and strictly increasing on r;

(A2) The tracking parameter δ satisfies the condition:

$$\frac{1}{2c} \sup_{r\in[r_{min}, r_{max}]} \frac{dp(r)}{dr} \le \delta \le \frac{a}{c} \inf_{r\in[r_{min}, r_{max}]} \frac{dp(r)}{dr},$$

for some a, with $\frac{1}{2} < a < 1$, where r_{min} and r_{max} denote the minimal and maximum value of r respectively.

In Theorem 3.2 of Ma et al. [2010] we showed that, under assumptions (A1) and (A2) and by implementing the update procedure (P1)-(P4) ahead of charging interval $[0, T]$, the infinite-population system converges to a unique Nash equilibrium $\mathbf{u}^*$. Moreover this unique Nash equilibrium satisfies the desired valley-filling property, which can be stated precisely as,

$$\overline{\mathbf{u}}_t^* \ge \overline{\mathbf{u}}_s^*, \quad \text{in case } d_{t|0} \le d_{s|0}, \tag{9a}$$

$$d_{t|0} + \overline{\mathbf{u}}_t^* \le d_{s|0} + \overline{\mathbf{u}}_s^*, \quad \text{in case } d_{t|0} \le d_{s|0}, \tag{9b}$$

$$d_{t|0} + \overline{\mathbf{u}}_t^* = B, \quad \text{for all } t \in \widehat{\mathcal{T}}, \tag{9c}$$

for some constant B, and where $\widehat{\mathcal{T}}$ represents a collection of instants t where $u_{nt}^* > 0$ for all n. The proof of this result is given as Theorem 3.3 in Ma et al. [2010].

Note that the Nash equilibrium properties (9) do not guarantee perfect globally optimal valley filling because there may be intervals in which not all PEVs charge. In this case B cannot be expressed in closed form. Nevertheless in the case of homogeneous PEV populations, the identical individual agent strategies $\mathbf{u}_n^*$ are coincident with the average charging strategy $\overline{\mathbf{u}}^*$. It follows that the Nash equilibrium properties (9) are equivalent to:

$$u_{nt}^* = \overline{\mathbf{u}}_t^* = \max\{0, B - d_{t|0}\}, \tag{10}$$

where $B > 0$ takes the value that ensures $\sum_{t=0}^{T-1} u_{nt}^* = \frac{\beta^n}{\alpha^n}(1 - x_{n0})$, i.e., all PEVs are fully charged at the end of the charging interval. In the charging control of homogeneous PEV populations, the Nash equilibrium becomes a purely valley-filling charging strategy, with respect to the non-PEV demand estimation $d_{.|0}$. Figure 1 provides an illustration.

Although the theory, as stated, applies to infinite populations, in reality any implementation of the control strategy works for a finite, though large, number of PEVs. For a large population, the average mass behavior of the population with the n-th PEV excluded, denoted $\overline{\mathbf{u}}_{-n}$, can be approximated by the common mass trajectory $\overline{\mathbf{u}}$. Hence for large population size, the iterative process (P1)-(P4) converges to an ε-Nash equilibrium, for some $\varepsilon > 0$. It follows that ε shrinks to zero as the population size N approaches infinity.

To illustrate the decentralized charging control result, we present a numerical example that explores the convergence properties and valley-filling performance of the decentralized charging control process. We use a PEV population of 10^7, which is approximately 30% of all the vehicles in the region of the USA covered by the Midwest Independent System Operator (MISO). We also assume the normalized fixed generation capacity c in the MISO region is about 10 kW.

We consider decentralized charging control for a homogeneous PEV population with common battery size equal to 10 kWh. It is further assumed that all PEVs have an initial SOC of 15%, identical charging efficiency α of 85%, and maximum charging rate of 3 kW. We consider a 12 hour charging interval from 8:00pm one day to 8:00am

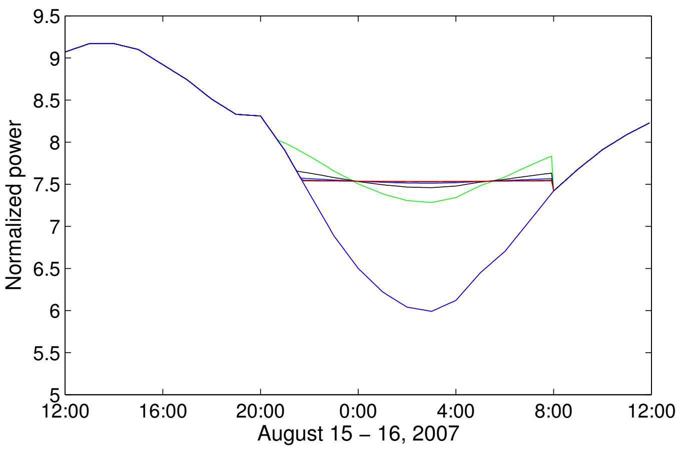

Fig. 1. Convergence to a Nash equilibrium for a homogeneous PEV population.

the following day. Each time step in (1) corresponds to 5 minutes.

We also consider a retail price function $p(r) = 0.15r^2$ and a tracking parameter $\delta = 0.015$. With these parameters, it is easy to verify that assumptions (A1) and (A2) are satisfied. Hence, it is guaranteed that the algorithm (P1)-(P4) will converge to a unique valley-filling Nash equilibrium. Figure 1 shows the total (PEV plus non-PEV) demand obtained at each iteration of the algorithm. Convergence to the valley-filling load pattern is obtained in a few iterations.

Unless otherwise specified, this example will form the basis for the later illustrations.

3. DECENTRALIZED MPC METHODS FOR PEV CHARGING CONTROLS

The decentralized charging control strategy in Section 2 computes optimal control trajectories with respect to non-PEV demand predicted at the beginning of the charging interval. Any subsequent disturbances or changes in the forecast are not dealt with by the approach. As a consequence, depending on the accuracy of the forecast at the beginning of the interval, the charging strategy $\mathbf{u}^*$ attained at convergence of (P1)-(P4) may actually result in suboptimal charging cost for each PEV.

In order to improve the robustness of the decentralized charging strategy, this section establishes a revised process that is consistent with model predictive control (MPC) methods. We first develop notation to support the MPC approach. At a time instant $t \in [0, T]$, all PEVs have access to a prediction over the subsequent interval $[t, T]$. We denote $u^n_{\cdot|t}$ as the *predicted charging strategy* of the n-th PEV over $[t, T]$, and $\mathbf{u}_{\cdot|t} \equiv (u^n_{\cdot|t}; n \in \mathcal{N})$ as the collection of charging strategies. The *predicted SOC trajectory* of the n-th PEV over $[t, T]$ is denoted $x^n_{\cdot|t}$. The evolution of $x^n_{\cdot|t}$ satisfies the transition equation (1) subject to a charging strategy $u^n_{\cdot|t}$. In the MPC formulation, this becomes

$$x^n_{s+1|t} = x^n_{s|t} + \frac{\alpha_n}{\beta_n} u^n_{s|t}, \quad \text{for each } s \in [t, T], \tag{11}$$

with $x^n_{t|t} \equiv x_{nt}$ being the actual measured SOC at time instant t. Similar to Section 2, $u^n_{\cdot|t}$ is an admissible full-charging control strategy for the n-th PEV over the prediction horizon $[t, T]$, if $u^n_{s|t} \geq 0$ for all $s \in [t, T]$, and $x^n_{T|t} = 1$. We define the set of predicted admissible full-charging strategies of the n-th PEV at time instant t as,

$$\mathcal{U}^t_n(x_{nt}) = \Big\{ u^n_{s|t}; u^n_{s|t} \geq 0, \sum_{s=t}^{T-1} u^n_{s|t} = \frac{\beta_n}{\alpha_n}(1 - x_{nt}) \Big\}. \tag{12}$$

As defined earlier, $d_{\cdot|t}$ is the prediction of non-PEV demand over the prediction horizon $[t, T]$, given information up to instant t. Also, $r_{\cdot|t} \triangleq \frac{1}{c}(d_{\cdot|t} + \overline{\mathbf{u}}_{\cdot|t})$ is the natural extension of (6)-(7).

Decentralized MPC charging control calls for each PEV of the infinite population to minimize its cost,

$$\min_{u^n_{\cdot|t} \in \mathcal{U}^t_n(x_{nt})} J_n(\mathbf{u}_{\cdot|t}) = \sum_{s=t}^{T-1} p(r_{s|t}) u^n_{s|t} + \delta (u^n_{s|t} - \overline{\mathbf{u}}_{s|t})^2, \tag{13}$$

where J_n is defined with respect to the demand forecast $d_{\cdot|t}$ and the collection of charging strategies $\mathbf{u}_{\cdot|t}$. The following algorithm describes a process for implementing this decentralized optimization:

(M1) At instant t, the utility generates a new non-PEV demand forecast $d_{\cdot|t}$ over the prediction interval $[t, T]$. The forecast and initial estimate of $\overline{\mathbf{u}}_{\cdot|t}$ is broadcast to all PEVs. This initial estimate is available from the converged solution of the MPC process at the previous interval, namely $\mathbf{u}^*_{\cdot|t-1}$.

(M2) Each PEV computes its optimal charging strategy $u^{n*}_{\cdot|t}$ over $[t, T]$ by minimizing its charging cost (13), which is defined with respect to its initial SOC state $x^n_{t|t}$, non-PEV demand $d_{\cdot|t}$, and the common average (or aggregate) predicted PEV charging strategy $\overline{\mathbf{u}}_{\cdot|t}$ broadcast by the utility.

(M3) The utility collects all individual charging strategies proposed in (M2), and updates the aggregate PEV demand. This updated aggregate PEV demand is rebroadcast to all PEVs.

(M4) Repeat (M2) and (M3) until the optimal strategies $(u^{n*}_{\cdot|t}; n \in \mathcal{N})$ proposed by the PEVs no longer change.

(M5) Each of the PEVs implements the charging control $u^{n*}_{t|t}$ at instant t as determined in (M4). The SOC subsequently evolves to its new value $x_{n,t+1}$ under the influence of this control input.

(M6) Set $t := t + 1$ and repeat the procedure (M1)-(M5) until the system reaches the end of the charging period at $t = T$.

A comparison between the optimal charging strategies of the open-loop process (P1)-(P2) and the MPC formulation (M1)-(M6) is provided by the following lemma.

Lemma 1. Suppose the prediction of non-PEV demand $d_{s|t}$ is coincident with $d_{s|0}$ for all $t \in [0, T]$ and $s \geq t$. Then the charging strategy $\mathbf{u}^*_{s|t}$ given by the MPC process (M1)-(M6) is identical to the Nash equilibrium $\mathbf{u}^*$ attainted by the procedure (P1)-(P4).

The Lemma follows directly from Bellman's principle of optimality, see for example Bertsekas [1995].

■

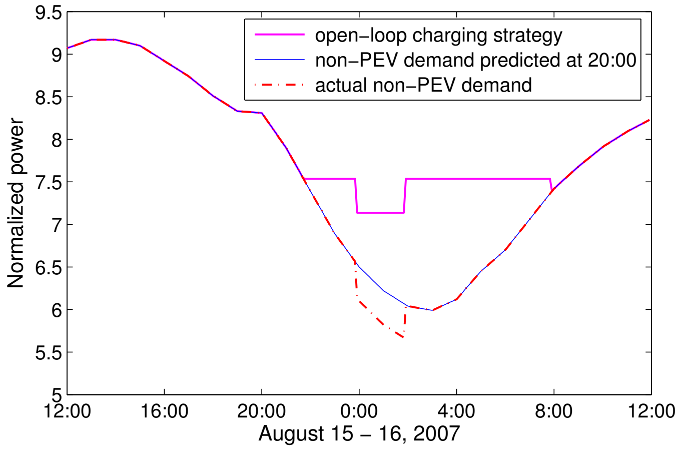

Fig. 2. Effect of a disturbance in non-PEV demand on open-loop charging.

4. ILLUSTRATING EXAMPLES

The example presented in Section 2 forms the basis for illustrating various aspects of the MPC process of Section 3. Firstly, the effects of disturbances on the open-loop strategy determined by (P1)-(P4) will be explored. Secondly, we will consider the influence of the demand forecast $d_{\cdot|t}$ on the control strategy provided by the MPC process (M1)-(M6).

The open-loop process determines a control strategy based on a forecast of non-PEV demand, and implements that strategy whether the actual demand follows that forecast or not. This is illustrated by comparing Figures 1 and 2. The control strategy is identical in both cases. In Figure 2, however, the actual demand undergoes a step reduction at 0:00 and recovers at 2:00. The control cannot respond, so PEV demand remains unchanged from its open-loop strategy. Hence total demand exactly follows the dip in non-PEV demand.

In Figure 3, non-PEV demand undergoes the same dip as in Figure 2. In this case, the MPC controller is active, so when the step reduction in non-PEV demand occurs, the controller makes a corresponding change to the forecast. However it has no knowledge of future deviations, so assumes the shape of the forecast over the remainder of the charging period will exactly match the original forecast, but minus the step reduction. Because it is the *relative* variation in demand across time intervals that influences the optimization (13) performed by each PEV, the MPC process will converge to a Nash equilibrium that is unchanged from its pre-disturbance shape. This argument again follows when the non-PEV demand steps back up at 2:00. Therefore if the forecast used by MPC has no updated knowledge of future changes in non-PEV demand, then the outcome exactly matches that of the open-loop strategy.

The load behavior shown in Figure 4 results when the updated MPC forecast takes into account changes in the relative shape of the non-PEV demand trajectory. Until 0:00, MPC operates with the same forecast as in previous cases. It follows that charging control over this period also matches the previous examples. At 0:00, when the

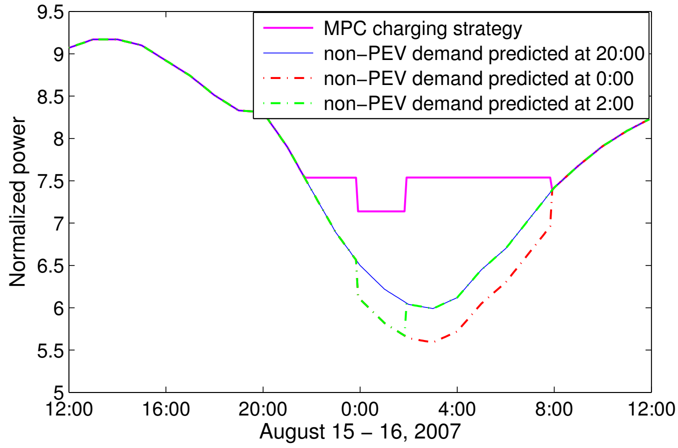

Fig. 3. MPC charging strategy with an imperfect prediction of non-PEV demand.

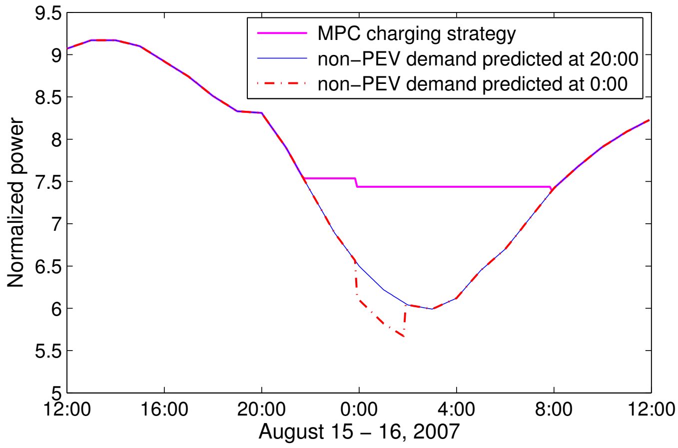

Fig. 4. MPC charging strategy with a perfect prediction of non-PEV demand.

step reduction in non-PEV demand occurs, the updated forecast includes both the reduction and subsequent step increase at 2:00. With this additional information, the PEV charging control strategy adjusts to again achieve valley filling. Figure 4 clearly shows this adjustment. Because valley filling is optimal, the total cost in this case is less than in the corresponding previous cases.

5. CONCLUSIONS

This paper extends the authors previous work on decentralized game-based PEV charging strategies by developing (and testing by simulation) a simple method to recompute optimal control strategies as updated forecasts are received. The strategy draws on concepts from model predictive control, most importantly the process of recomputing a new optimal control trajectory at each time step (and therefore dispensing with all but the first of the control values computed at the previous step). Interestingly, the "receding horizon" strategy that is broadly used in other MPC work is not applicable in this context because the control interval ends at a fixed time.

Numerical examples in this paper have demonstrated that if forecast updates are simply time-invariant shifts to the previous forecast, the optimal PEV charging strategy does not change. However, if differences between new and previous forecasts are time-dependent (and therefore the relative *shape* of the forecast quantities changes), the method described in this paper does modify the optimal PEV charging pattern. This suggests that forecast modifications must have long-term structure for the procedure to be most useful.

We have not formally examined the number of iterations required to reach an optimal trajectory at each time step. This must eventually be better understood, since the number of interactions, in combination with the capacity of the communications network and local computing resources, will determine how often forecast updates could be incorporated into the framework proposed in this paper.

Another important opportunity that could be incorporated into this load-control framework concerns the magnitude of ramps in total (PEV plus non-PEV) demand. In principle, the PEV population is not bound by ramp-rate constraints that limit the response of other devices, principally generators, that provide reserves and frequency regulation services. This allows PEVs to match very fast changes in demand or supply, alleviating the burden on supply-side options.

REFERENCES

D.P. Bertsekas. *Dynamic programming and optimal control*, volume I. Athena Scientific, 1995.

P. Denholm and W. Short. An evaluation of utility system impacts and benefits of optimally dispatched plug-in hybrid electric vehicles. Technical Report NREL/TP-620-40293, National Renewable Energy Labortary, October 2006.

F. Koyanagi and Y. Uriu. Modeling power consumption by electric vehicles and its impact on power demand. *Electrical Engineering in Japan*, 120(4):40–47, 1997.

F. Koyanagi, T. Inuzuka, Y. Uriu, and R. Yokoyama. Monte Carlo simulation on the demand impact by quick chargers for electric vehicles. In *Proceedings of the IEEE Power Engineering Society Summer Meeting*, volume 2, pages 1031–1036, Edmonton, Alberta, July 1999.

Z. Ma and D. Callaway. On the optimal charging control of large populations of plug-in electric vehicles. Technical report, Energy and Resources Group, University of California, Berkeley, 2011.

Z. Ma, D. Callaway, and I. Hiskens. Decentralized charging control for large populations of plug-in electric vehicles. In *Proceedings of the 49th IEEE Int. Conf. Decision and Control*, pages 206–212, Atlanta, GA, Dec. 2010.

Z. Ma, D. Callaway, and I. Hiskens. Decentralized charging control of large populations of plug-in electric vehicles. *IEEE Transactions on Control Systems Technology*, 2011. Accepted for publication.

S. Rahman and G.B. Shrestha. An investigation into the impact of electric vehicle load on the electric utility distribution system. *IEEE Transactions on Power Delivery*, 8(2):591–597, 1993.