2010 IEEE International Conference on Control Applications
Part of 2010 IEEE Multi-Conference on Systems and Control
Yokohama, Japan, September 8-10, 2010

Decentralized Charging Control for Large Populations of Plug-in Electric Vehicles: Application of the Nash Certainty Equivalence Principle

Zhongjing Ma Duncan Callaway Ian Hiskens

Abstract— **The paper develops and illustrates a novel decentralized charging control algorithm for large populations of plug-in electric vehicles (PEVs). The proposed algorithm is an application of the so-called Nash certainty equivalence principle (or mean-field games.) The control scheme seeks to achieve social optimality by establishing a PEV charging schedule that fills the overnight demand valley. The paper discusses implementation issues and computational complexity, and illustrates concepts with various numerical examples.**

Keywords: **Decentralized control; plug-in electric vehicles (PEVs); Nash certainty equivalence (NCE) principle; 'valley-fill' charging control.**

I. Introduction

In order to reduce the emission of green-house gases and reliance on exhaustible petroleum sources, high penetrations of plug-in electric vehicles are expected to substitute the current conventional petroleum-combustion vehicles over the next few decades. The electricity demand from this large populations of PEVs may have a significant impact on the electrical power grid. For example, suppose that 30% of the 234 million conventional vehicles in the US were substituted by PEVs, that the average size of the PEV batteries is about 10 kWh, and that the charging rate of each PEV is about 2 kW. The total charging load would be 140 GW, which is 18% of the US summer peak load of 780 GW.

Quite a few studies have been undertaken recently to explore the potential impacts of high penetrations of PEVs on the power grid [1], [2], [3], [4]. In [5], we study centralized optimal charging control, for large populations of homogeneous PEVs. This work shows that under certain reasonable conditions, the control strategy results in *valley filling*, i.e. the total demand, consisting of aggregated PEV charging load and non-PEV demand, is constant during charging intervals. Note that all these proposals assume that the utility can directly control the charging rates of individual PEVs.

The implementation of centralized charging control for large populations of PEVs is computationally intractable in general. It may also be impractical, due to the possible reluctance of PEV owners to allow their utility to directly control vehicle charging rates. In this paper, we suppose that each of the PEV agents implements local charging controls for its own PEV.

The (electricity) charging price, seen by all PEVs, is responsive to the total demand of the grid, which is the summation of the inelastic non-PEV base demand together with the aggregated charging rates of the whole population of PEVs. Because of the coupling through this common price signal, each PEV agent effectively interacts with the average charging strategy of the rest of the PEV population. As the population grows substantially, the influence of each individual PEV on that average charging strategy becomes negligible. Accordingly, for large populations, the average charging strategy seen by every PEV is identical.

As a consequence, considering the charging controls for an infinite PEV population, a collection of local charging controls is a Nash equilibrium (NE), if

(i) Each of the local controls is optimal with respect to one commonly observed charging trajectory, and
(ii) The average of these local optimal charging controls is equal to the common trajectory, i.e. the average charging strategy is collectively reproduced by the local optimal control laws.

This result is referred to as the Nash certainty equivalence (NCE) principle, as proposed by Huang *et al.* [6], [7]. This framework has connections with mean-field game models that were studied by Lasry and Lions [8], [9], and close connections with the notion of oblivious equilibrium proposed by Weintraub, Benkard, and Van Roy [10] via a mean-field approximation.

Under certain reasonable conditions, the paper demonstrates via illustrative examples that there exists a Nash equilibrium. Moreover assuming that the electricity price is strictly increasing with respect to the total demand, it is verified that at a Nash equilibrium, the aggregated charging rates (of the PEV population) almost achieve 'valley-fill' (hence are nearly socially optimal.) For homogeneous PEVs, the charging control turns out to be a 'valley-fill' strategy.

The paper is organized as follows. A class of PEV decentralized charging control problems is formulated in Section II. Section III introduces the so-called Nash certainty equivalence (NCE) principle. Using the NCE methodology, an algorithm is designed to implement the underlying decentralized control strategy. The illustrative examples of Section IV demonstrate algorithm behaviour. The cost performance of the decentralized charging strategy is compared for different system specifications. Section V proposes an approach to handling unpredicted changes in system conditions. Conclusions are presented in Section VI, along with a

Zhongjing Ma is with the Center for Sustainable Systems (CSS) and the School of Natural Resources and Environment, University of Michigan, e-mail: mazhong@umich.edu.

Duncan Callaway is with the Energy and Resources Group, University of California, Berkeley, e-mail: dcal@berkeley.edu.

Ian Hiskens is with the Department of Electrical Engineering and Computer Science, University of Michigan, email: hiskens@umich.edu.

Research supported by the Michigan Public Service Commission through grant PSC-08-20, and the National Science Foundation through EFRI-RESIN grant 0835995.

TABLE I
LIST OF SYMBOLS (FOR PEV n).

$x_t^n \in [0\ 1]$	state of charge (SOC) at time t
x_0^n	initial SOC
$u_t^n \geq 0$	charging rate at time t
α^n	charging efficiency
β^n	battery size
T_0	starting charging time (identical for all PEVs)
T	terminal charging time (identical for all PEVs)

discussion of future research directions.

II. DECENTRALIZED CHARGING CONTROL PROBLEMS FOR LARGE POPULATIONS OF PEVS

We consider charging control for a significant penetration of PEVs with a population size equal to N. For individual PEV n, we adopt the notation of Table I. It follows that the charging dynamics can be written,

$$x_{t+1}^n = x_t^n + \frac{\alpha^n}{\beta^n} u_t^n, \qquad t = T_0, \cdots, T-1 \tag{1}$$

with an initial state-of-charge (SOC) of x_0^n. We assume the PEV is fully charged at the end of the charging interval, $x_T^n = 1$. Accordingly, we define the *set of feasible full charging controls*,

$$\mathcal{U}^n \triangleq \Big\{ u^n \equiv \big(u_{T_0}^n, \cdots, u_{T-1}^n\big);\ \text{s.t. } u_t^n \geq 0,\ x_T^n = 1 \Big\}.$$

We denote $\mathbf{u} \triangleq \big\{u^n; 1 \leq n \leq N\big\}$ as the collection of charging rates of the whole PEV population, and $\mathbf{u}^{-n} \triangleq \big\{u^m; m \neq n\big\}$ as the collection of charging rates for the PEV population excluding the n-th PEV.

Coordinated control of PEV charging often assumes a centralized control framework, where the utility controls the charging rates for all PEVs. The objective is to implement a collection of PEV charging rates that achieve the dual objectives, 1) the aggregated PEV load fills the over-night valley of the base load, and 2) every PEV is fully charged at the end of its charging interval.

In contrast, this paper proposes a charging control scenario where individual PEVs minimize their own operating cost by implementing a charging strategy that takes into account the collection of charging strategies adopted by other agents. More specifically, consider a collection of charging strategies $\mathbf{u}$, with each $u^n \in \mathcal{U}^n$, and suppose that the cost function of agent n is given by

$$J^n(\mathbf{u}) \triangleq \sum_{t=T_0}^{T-1} \Big\{ p_t u_t^n + \delta\big(u_t^n - \text{avg}(\mathbf{u}_t)\big)^2 \Big\} \tag{2}$$

where δ is a non-negative constant, and $\text{avg}(\mathbf{u}_t) = \frac{1}{N}\sum_{n=1}^{N} u_t^n$. The price $p_t \equiv p(d_t + N\text{avg}(\mathbf{u}_t))$ denotes the electricity price at instant t, which is dependent upon the inelastic demand d_t and the total PEV demand. It follows from (2) that each agent's optimal charging strategy must achieve a trade-off between the total electricity cost pu^n and the cost incurred in deviating from the average behaviour of the PEV population $(u^n - \text{avg}(\mathbf{u}))^2$.

The tracking cost term may be thought of as a real-time regulation fee required by the utility for PEV agents to join the decentralized control scheme. The examples in Section IV will illustrate that the small tracking costs are more than compensated by cost savings that arise from valley filling.

The underlying decentralized PEV charging control scheme is a *finite-horizon noncooperative dynamic game*. Each PEV agent shares with other PEV agents the (limited, valuable and divisible) electricity resources beyond the inelastic non-PEV demand base d, and also tracks the average charging strategy of the whole population. A collection of charging controls $\mathbf{u}$ is a Nash equilibrium if, for all n, $u^n \in \mathcal{U}^n$ is a charging strategy that minimizes the operation cost (2) with respect to $\mathbf{u}^{-n}$.

So far we have formulated the decentralized charging control problem as a class of dynamic games. Such problems are, however, generally computationally intractable for large population size N. This issue is addressed in Section III through the introduction of the Nash certainty equivalence (NCE) principle, and the use of NCE in the design of an algorithm to implement this class of problem. The illustrative examples of Section IV demonstrate the implementation of the decentralized charging control algorithm. The cost performance of the decentralized charging strategy under different system specifications will be discussed.

III. IMPLEMENTATION OF DECENTRALIZED PEV CHARGING CONTROL USING THE NCE METHODOLOGY

The Nash certainty equivalence principle was first explored by M. Huang *et al.* [6], [11], [7] in the context of large scale dynamic games for sets of weakly coupled stochastic control systems. In such systems, the dynamics and the payoff/cost function of each individual agent are influenced by certain average values of the collection of agents.

At an established Nash equilibrium, each agent reacts optimally with respect to its local state and the collectively generated average trajectory of all other agents. These trajectories are approximated by an identical deterministic infinite population limit (associated with the mean field or ensemble statistics of the random agents) which is the solution of a particular fixed point problem.

It is straightforward to check that the decentralized charging control problem formulated in Section II, with the SOC dynamics given by (1) and the cost functions of individual PEV agents specified by (2), is in the framework of the NCE principle. It follows that an infinite collection of charging strategies $\{u^n; 1 \leq n < \infty\}$ is a Nash equilibrium, if u^n is an optimal feedback charging strategy with respect to a common $\overline{u}$ given by the average value of the other agents' behaviours. In other words, $\overline{u} = \text{avg}(\mathbf{u}^{-n})$ for all n.

For any finite population size N, however, the NCE methodology gives the weaker result that $\{u^n; 1 \leq n \leq N\}$ is an ε-NE, for some positive ε. All the $\text{avg}(\mathbf{u}^{-n})$ are only approximately equal in this case, because of the finite population size N. Nevertheless the value of the error ε tends to zero as the population size N grows.

Implementation of the NCE-based decentralized control is achieved through a *charging negotiation procedure*, which takes place at some time prior to the actual charging interval:

(S1) The utility broadcasts the prediction of non-PEV base demand d to all the PEV agents.

(S2) Each of the PEVs proposes a charging control minimizing its charging cost with respect to a common aggregate PEV demand broadcast by the utility.

(S3) The utility collects all the individual optimal charging strategies proposed in (S2), and updates the aggregate PEV demand corresponding to the proposed charging strategies. This updated aggregate PEV demand is rebroadcast to all of the PEVs.

(S4) Repeat (S2) and (S3) until the optimal strategies proposed by the agents no longer change.

At convergence, the collection of proposed individual charging strategies is an NE. Some time later, when the charging start time T_0 is reached, each PEV implements the optimal strategy obtained from (S1)-(S4).

In the negotiation procedure (S1)-(S4), each of the PEV agents independently updates its own optimal feedback charging strategy with respect to the one-dimensional average value of all PEV agent strategies. The computational complexity of the underlying decentralized control strategy is therefore independent of the PEV population size N.

IV. Performance Investigations

This section demonstrates, through a number of illustrative examples, the convergence properties and social optimality (or valley-filling) performance of the NCE-based decentralized charging control process. A range of system conditions are considered, with a particular focus on the tracking-cost parameter δ. Rigorous proofs of convergence and optimality are presented in [12]. This paper focuses on the concepts rather than the details.

It is common to assume that the electricity retail price p is strictly increasing with total demand. This assumption is consistent with the aggregated marginal cost curve in a deregulated electricity market. For example, the cloud of green dots in Figure 1 shows the day-ahead marginal prices for electricity during the summers of 2007 and 2008 for the market managed by the Midwest Independent System Operator (MISO). The blue curve provides an estimate of marginal cost curve obtained through "kernel nonparametric regression (NPR)." For simplicity, the following illustrative examples make use of the smoothed electricity retail price curve shown by the black curve in Figure 1. Note that this curve is strictly increasing.

All the numerical examples assume the total generation capacity in the MISO region is about 10.5×10^4 MW, which corresponds to the maximum load supplied in the MISO region during 2007 and 2008. A PEV population of 10^7 will be used. That number of vehicles is roughly equivalent to 30% of all the vehicles in the MISO region. It is further assumed that all PEVs have an initial SOC of 15%, and

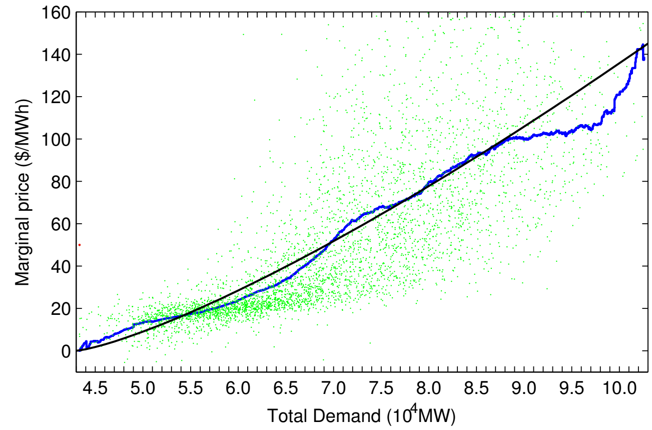

Fig. 1. Day-ahead electricity marginal prices in summer, 2007 and 2008 by quantity demanded in Midwest ISO, an associated "kernel nonparametric regression (NPR) estimate" and a strictly increasing curve approximate to NPR estimate marginal price.

identical charging efficiency α of 85%. We consider a 24-hour charging interval during the summer of 2007, from noon on August 6 to noon on August 7. Other parameters, such as PEV battery size β^n and the tracking-cost parameter δ, will be specified for each of the examples.

The tracking-cost parameter δ plays a particularly important role in the convergence behaviour of the NCE-based control strategy. The following examples explore its impact. Initially the tracking cost will be set to zero. This is achieved by setting $\delta = 0$, giving

$$J^n(\mathbf{u}) = \sum_{t=T_0}^{T-1} p_t u_t^n, \quad \text{with } p_t \equiv p(d_t + N\text{avg}(\mathbf{u}_t)).$$

Recall from Figure 1 that price is an increasing function of total demand. Each PEV will seek to charge when the price is low, which corresponds to forecast demand being low. Accordingly, the maximum charging rate will occur when total demand is at its minimum. Notice though that PEV agents are only aware of forecast demand, and are oblivious to each others actions. This results in all agents acting in unison. Consequently, the iterative procedure (S1)-(S4) oscillates indefinitely.

This behaviour is illustrated in Figure 2. The figure shows the outcome for a homogeneous population of PEVs that all have identical battery size of 10 kWh and maximum charging rate of 3 kW. Based on the forecast non-PEV demand, all agents choose the same optimal charging strategy. This results in the demand spike shown by the central blue curve. At the next iteration of the negotiation procedure (S1)-(S4), all PEVs see the blue demand curve and respond accordingly. The double peaks shown in green are the result. Subsequent iterations of (S1)-(S4) oscillate between the blue and green curves.

It may be concluded that by adopting a zero or small tracking-cost parameter δ, all agents seek to utilize the cheapest electricity resources. Consequently, the cheap resources during the off-peak period are over utilized, creating a new

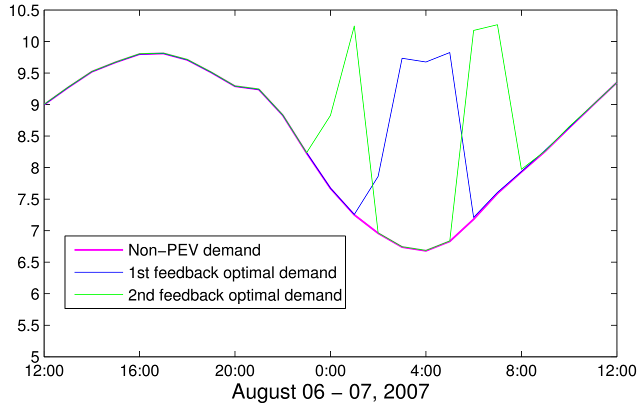

Fig. 2. Optimal charging strategy for homogeneous PEVs with zero tracking cost, battery size 10 kWh, and maximum charging rate 3 kW.

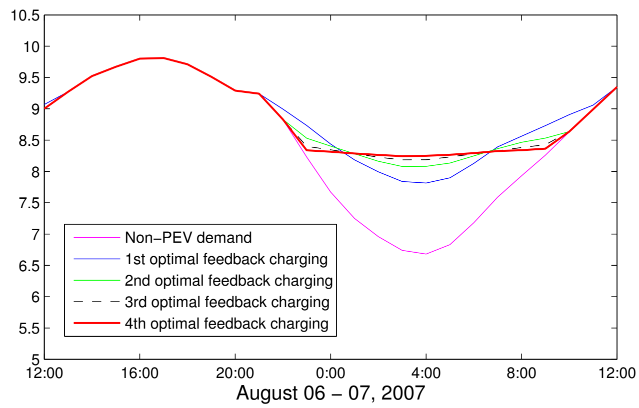

Fig. 3. Converging to a Nash equilibrium ('red curve') for homogeneous PEVs with zero initial average strategy and tracking-cost parameter $\delta = 0.007$.

peak. That new peak is avoided at the next negotiation cycle. Because of this cycling, the system cannot converge to a Nash equilibrium.

By choosing a sufficiently large tracking-cost parameter δ, however, the best charging strategy of every agent will take into account the trade-off between the charging price and the cost of deviating from the average charging strategy. As a result, agents are less aggressive in pursuing the cheaper electricity resources. Figure 3 displays the iterates of the optimal charging strategies for a tracking-cost parameter $\delta = 0.007$. In this case, the charging system converges in four negotiation cycles to the Nash equilibrium, which corresponds to valley-filling.

In [12], we quantify a range for tracking-cost parameter δ that ensures the negotiation procedure (S1)-(S4) converges to the unique Nash equilibrium. Moreover, this range is unrelated to the initial choice for the average charging strategy. As an illustrative example, consider the unusual initial average strategy shown by the solid black curve in Figure 4. The sequence of curves in Figure 4 indicate that the system converges, in a few negotiation cycles, to the same 'valley-filling' Nash equilibrium as displayed in Figure 3.

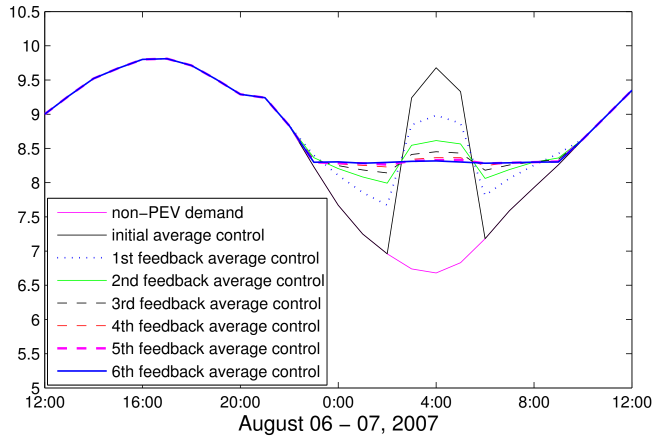

Fig. 4. Converging to a Nash equilibrium ('blue curve') for homogeneous PEVs with non-zero initial average strategy and tracking-cost parameter $\delta = 0.007$.

TABLE II

SPECIFICATIONS FOR TWO CLASSES OF PEV POPULATIONS

	PEV population 1	PEV population 2
Population size	$N^1 = 0.5 \times 10^7$	$N^2 = 0.5 \times 10^7$
Battery size	$\beta^1 = 10$	$\beta^2 = 20$

Up to now, the focus has been on homogeneous populations. However the convergence guarantee can be extended to the more general case of heterogeneous populations. Figure 5 displays the converged Nash equilibrium for decentralized charging of two classes of populations, whose specifications are given in Table II. The tracking-cost parameter $\delta = 0.007$ was again chosen. The associated Nash equilibrium, shown by the solid red line in Figure 5, gives a charging strategy that is almost valley filling.

In summary, this section has demonstrated that by introducing a certain positive tracking cost, the decentralized charging control scheme converges to a unique Nash equilibrium, in a finite number of negotiation cycles. The converged solution is 'valley-filling' in the case of homogeneous PEV populations, and nearly 'valley-filling' for heterogeneous PEV populations. Furthermore, for homogeneous decentralized charging control problems, the tracking costs are transitional, occurring only during the negotiation procedure. At the 'valley-filling' NE, the tracking cost equals zero, since each agent's charging strategy coincides with the average over all strategies. For heterogeneous control problems, the tracking cost at convergence is positive, since each agent's charging strategy is generally not the same as the average across all strategies. Note though that the tracking cost is small, since the tracking-cost parameter is much smaller than the electricity price.

V. ADAPTING TO CHANGING SYSTEM CONDITIONS

The negotiation procedure (S1)-(S4) occurs ahead of the actual charging interval, and is based on predictions of non-PEV demand and of the PEV population to be charged. Conditions invariably differ during the actual charging in-

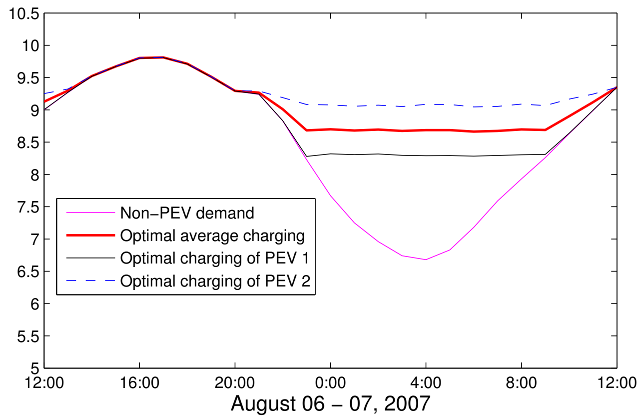

Fig. 5. A converged Nash equilibrium for the two classes of PEV populations specified in Table II, and with $\delta = 0.007$

terval. The non-PEV demand will never exactly match the prediction, and there will always be some mismatch between the PEVs that participated in the prior negotiations and those that actually charge. Also, major disturbances such as loss of a large generator are always a possibility.

The decentralized control algorithm can be reformulated to allow tracking of system variations. This tracking process effectively becomes a form of model predictive control [13]. Real-time information is collected, and a revised prediction of demand is broadcast to all PEV agents. Negotiations (S2)-(S4) proceed to determine the optimal updated schedule for each PEV. The PEVs then charge according to that new schedule.

VI. Conclusions and Future Research

The paper introduces a novel decentralized charging control algorithm for large populations of PEVs. It describes an application of the Nash certainty equivalence (NCE) methodology, and provides an overview of the more rigorous investigations reported in [12].

It is shown in the paper that if PEVs act to minimize their cost, with no regard to other PEVs, the iterative scheduling process is unlikely to be convergent. However, by penalizing PEVs for deviating from the average behaviour of all other PEVs, the scheduling process is guaranteed to converge to a unique Nash equilibrium. If the PEV population is homogeneous, the Nash equilibrium will optimally fill the overnight demand valley. Furthermore, all PEVs will adopt identical charging strategies, and the cost penalty associated with deviating from the average will be zero. On the other hand, if the PEV population is heterogeneous, PEV charging strategies will be similar, but not identical, and the cost penalty for deviating from average will be small but not zero.

Future research will explore a number of new directions. Current work assumes that the non-PEV demand is deterministic and predictable. We will generalize this work to incorporate stochastic non-PEV demand forecasting models, and will evaluate trade-offs between robustness and conservativeness. Large populations of PEVs are conceptually similar to collections of distributed energy resources (DERs). It is interesting to consider the extension of the PEV control scheme to incorporate high penetration of stochastic renewable energy sources, such as wind and solar. Along a similar line, PEVs may ultimately have “vehicle-to-grid” (V2G) capability, whereby they could acts as loads at certain times and sources at other times. These possibilities will be incorporated into the NCE-based decentralized control structure.

References

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

[2] S. Rahman and G. Shrestha, “An investigation into the impact of electric vehicle load on the electric utility distribution system,” *IEEE Transactions on Power Delivery*, vol. 8, no. 2, pp. 591–597, 1993.

[3] F. Koyanagi and Y. Uriu, “Modeling power consumption by electric vehicles and its impact on power demand,” *Electrical Engineering in Japan*, vol. 120, no. 4, pp. 40–47, 1997.

[4] 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*, vol. 2, 18-22 July 1999, pp. 1031–1036.

[5] Z. Ma and D. Callaway, “Centralized optimal charging of large population plug-in electric vehicles,” *Technical Report*, June 2010.

[6] M. Huang, P. Caines, and R. Malhamé, “Individual and mass behaviour in large population stochastic wireless power control problems: centralized and nash equilibrium solutions,” in *42th IEEE Int. Conf. Decision and Control*, Maui, Hawaii, December 2003, pp. 98–103.

[7] ——, “Large-population cost-coupled LQG problems with non-uniform agents: individual-mass behaviour and decentralized epsilon-nash equilibria,” *IEEE Transactions on Automatic Control*, vol. 52, no. 9, pp. 1560–1571, 2007.

[8] J.-M. Lasry and P.-L. Lions, “Jeux à champ moyen. I – le cas stationnaire,” *C. R. Acad. Sci. Paris, Ser. I*, vol. 343, pp. 619–625, 2006.

[9] ——, “Mean field games,” *Japan. J. Math.*, no. 2, pp. 229–260, 2007.

[10] G. Weintraub, C. Benkard, and B. V. Roy, “Markov perfect industry dynamics with many firm,” *Econometrica*, vol. 76, no. 6, pp. 1375–1411, November 2008.

[11] M. Huang, R. Malhamé, and P. Caines, “Nash certainty equivalence in large population stochastic dynamic games: connections with the physics of interacting particle systems,” in *45th IEEE Int. Conf. Decision and Control*, San Diego, CA, December 2006, pp. 4921–4926.

[12] Z. Ma, D. Callaway, and I. Hiskens, “Decentralized charging control for large population plug-in electric vehicles: Nash certainty equivalence principle,” *Technical Report*, June 2010.

[13] E. Camacho and C. Bordons, *Model Predictive Control*, 2nd ed. Berlin: Springer, 2004.