



Lawrence Berkeley National Laboratory

Real-Time Charging Strategies for an Electric Vehicle Aggregator to Provide Ancillary Services

George Wenzel¹, Matias Negrete-Pincetic¹, Daniel Olivares¹, Jason MacDonald², and Duncan S. Callaway³

¹Pontificia Universidad Católica de Chile

²Lawrence Berkeley National Laboratory

³University of California, Berkeley

March 2017


Published in *IEEE Transactions on Smart Grid*

DOI: 10.1109/TSG.2017.2681961

Please cite as: G. Wenzel; M. Negrete-Pincetic; D. E. Olivares; J. MacDonald; D. S. Callaway, "Real-Time Charging Strategies for an Electric Vehicle Aggregator to Provide Ancillary Services," in *IEEE Transactions on Smart Grid*, vol.PP, no.99, pp.1-1

Partial funding for this work by the California Energy Commission and the U.S. Department of Defense's Environmental Security Technology Certification Program (ESTCP).

© 2017 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.



Lawrence Berkeley National Laboratory

One Cyclotron Road | Berkeley, California 94720

Disclaimer

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Real-Time Charging Strategies for an Electric Vehicle Aggregator to Provide Ancillary Services

George Wenzel, Matias Negrete-Pincetic, Daniel Olivares, Jason MacDonald, and Duncan S. Callaway

Abstract—Real-time charging strategies, in the context of vehicle to grid (V2G) technology, are needed to enable the use of electric vehicle (EV) fleets batteries to provide ancillary services (AS). In this paper we develop tools to manage charging and discharging in a fleet to track an Automatic Generation Control (AGC) signal when aggregated. We propose a real-time controller that considers bidirectional charging efficiency and extend it to study the effect of looking ahead when implementing Model Predictive Control (MPC). Simulations show that the controller improves tracking error as compared with benchmark scheduling algorithms, as well as regulation capacity and battery cycling.

Index Terms—Electric Vehicles, Resource Scheduling, Ancillary Services, Vehicle to Grid

I. INTRODUCTION

NEW generation, demand, transmission and storage systems are presenting opportunities to increase power system flexibility. To capitalize on these opportunities, algorithms that coordinate distributed resources need to manage charging cost, efficiency and energy and power constraints. This paper focuses on the potential of a subset of flexible storage technologies, specifically electric vehicles.

Increased variability in power generation due to renewable energy integration makes storage capacity particularly valuable [1]. However stationary batteries are currently too expensive for most grid-tied applications despite their decreasing cost [2], [3]. Electric vehicle (EV) batteries can be used during their idle time when parked to extract/inject power from/to the grid in the same way that stationary batteries might. By creating revenue for EV owners – and lowering the total cost of EV ownership – this vehicle to grid (V2G) framework could provide a cost-effective means to add storage capacity to the grid.

Large shares of renewable generation are being integrated into the power grid mainly due to environmental concerns and energy supply issues. However, the key characteristics of renewable resources in terms of volatility, intermittency and uncertainty present great operational challenges for power systems. In particular, power system flexibility, defined as the capacity to respond to changes in load and generation, becomes critical for systems with large penetration of volatile resources [4]. Some authors show that the value of resources

with volatility above a certain threshold can be diminished without proper flexibility capabilities [5], [6].

New ways to obtain the required flexibility are under development. According to [7], flexibility is present in generation (ramping capability), transmission (bottlenecks and access), demand (demand response, storage and load control) and system operation (institutional factors, information and real-time decisions). The idea of exploiting the flexibility associated with the demand side has been widely investigated, and markets aiming for flexible loads to be serviced by zero-marginal cost renewable generation have been designed [8], [9].

One possible way of adding flexibility to power systems is through the V2G concept, for which several implementation projects and impact studies have been reported such as [10], [11]. In the literature, the concept [12], [13] and impacts [14], [15] of using EVs for grid stabilization have been extensively investigated. Numerous charging strategies to coordinate the response of EV fleets to provide frequency regulation services have been developed in the recent years [16]–[27]. These strategies can be separated into two main groups: centralized strategies, which use an EV aggregator as a middleman between the ISO and the EVs (e.g.: [23]–[27]) and decentralized strategies, which do not use an aggregator to coordinate individual EV charging commands (e.g.: [16]–[22]). While decentralized strategies preserve individual authority over charging schedules, centralized strategies allow to reduce the uncertainty over total available power and energy, which simplifies the interaction between the electricity market and each individual EV as shown in [28]. Furthermore, aggregation is still necessary to fulfill minimum power capacity requirements imposed by the ISO to participate in the ancillary services (AS) market.

In general, centralized strategies provide frequency regulation services by tracking regulation signals sent by the ISO; however, it is also possible to follow other objectives, such as reducing the Area Control Error (ACE) [24]. These regulation signals may require both extracting and providing power from/to the EV fleets, but only some papers [23], [26] consider bidirectional charging in their models. From these works, only [26] considers charging and discharging efficiencies, applied to a discrete set of possible charging rates (charging, idle, discharging). On the other hand, most of the works dealing with decentralized charging strategies, (e.g., [17], [19], [21], [22]) take into account bidirectional charging; however, only a few consider bidirectional efficiencies [17], [19]. Nevertheless, the latter references do not optimize their charging commands in response to market signals.

GW, MNP and DO are with the OCM-Lab at the Department of Electrical Engineering at Pontificia Universidad Católica de Chile, Santiago 7820436, Chile (e-mail: gawenzel@uc.cl). DC is with the Energy Resources Group at the University of California, Berkeley. JM is with the Lawrence Berkeley National Lab.

This research was partially funded by the Complex Engineering Systems Institute, ISCI (ICM-FIC: P05-004-F, CONICYT: FB0816), the Fondecyt Grant 11140621 and US NSF Grant CNS-1239467.

Due to the short period of regulation signals (4 seconds in CAISO) EV aggregators require fast strategies that enable them to distribute charging/discharging commands among EVs in real-time, while achieving the best possible performance [29]. This has led to the consideration of heuristic algorithms such as Earliest Deadline First (EDF) or Least Laxity First (LLF) for this purpose. However, these heuristics can negatively impact the performance of the batteries, especially when compared to alternative algorithms based on convex optimization formulations, as discussed in [30].

Model Predictive Control (MPC) is an optimization-based control method that can be used to track signals for which forecasts are available in real-time [31]. In particular, it can be used by an EV aggregator to track regulation signals, as shown in [26], where the authors propose to schedule charging/discharging commands based on a Linear Quadratic Regulator (LQR) that tracks the regulation signal, with only three possible charging rates.

A practical implementation of the V2G concept is being demonstrated on an operational fleet at the LAAFB. This project uses a hierarchical control framework, in which day-ahead and hour-ahead electricity market participation and charging schedule are handled by an optimization platform: DER-CAM (Distributed Energy Resources Customer Adoption Model) developed at Lawrence Berkeley National Laboratory (LBNL). DER-CAM optimizes distributed energy resources operation over economic and environmental objectives [32], [33]. DER-CAM performs a constrained economic optimization to generate bids for bulk energy and AS markets based on the forecasts of vehicle usage by calculating the vehicles' state of charge (SoC). However, the time resolution of DER-CAM's optimization is not suitable for responding to uncertain Automatic Generation Control (AGC) signals within a few seconds, which is key for achieving an accurate response to such signals. Scheduling methods must be designed to allow real-time operation of the fleet, and these methods must distribute power among the vehicles while following an uncertain AGC signal. Figure 1 depicts a schematic diagram of the control hierarchy for the LAAFB V2G project, showing the interaction of the real-time distribution developed for that project with the rest of the project. In previous work [30], a real-time controller based on convex optimization was described, and it was shown that better results can be achieved with that controller as compared to some benchmarks. This control algorithm is currently integrated into the EV fleet management platform developed by Kisensum, Inc. [34] for the project.

This paper develops a framework for designing real-time charging controllers to operate an EVs fleet participating in the AS market. We extend and refine an earlier conference paper [30] to improve regulation capacity and accuracy in following AGC signals¹ as compared with simpler approaches. A set of different controllers is designed and tested, from which the more complex approaches achieve better performance (accuracy) and less cycling. In specific, the contributions include:

¹Accuracy is particularly relevant because of performance payments.

TABLE I
PARAMETER DESCRIPTION FOR TASK i

Parameter	Description
r_{ik}	Nominal trajectory for time step k
E_{ik}	State of Charge for time step k
β_i^+, β_i^-	State of Charge physical limits
E_i^+, E_i^-	State of Charge goal at departure
a_i, d_i	Arrival/departure time
p_{ik}	Charging/discharging rate for time step k
m_i^+, m_i^-	Charging/discharging rate limits
η_i^+, η_i^-	Charging/discharging efficiency
ϕ_{ik}	Laxity for time step k

- The development of a modelling framework capable of handling bidirectional charging resources with efficiency considerations.
- The design and assessment of several real-time controllers with different levels of complexity (myopic, heuristic based and predictive) and features (e.g., reduction of cycling behaviour).
- The thorough performance assessment of the proposed models in comparison with non-predictive benchmarks, via extensive simulations under a range of scenarios, for both the regulation signal and the efficiency of the batteries.

This paper is organized as follows. Section II describes the models used for the batteries, the task concept and market participation. Section III presents the proposed controller and the benchmarks. Section IV describes an MPC version of the controller. Simulation results are presented and discussed in Section V, and Section VI explains this work's main conclusions.

II. PROBLEM SETUP

In the proposed framework, the functional unit is the battery, and it must be characterized in terms of a set of parameters. The task concept is used to describe the characteristics of a battery when it is active, and thus can be used within its physical limits by the EV aggregator to provide frequency regulation services, as opposed to an inactive task. The laxity concept is used as well to describe each task's flexibility [35].

A. Tasks and Batteries

An aggregator coordinating a fleet of EVs in real time faces, at each time step $k \in \{1, \dots, T\}$, the challenge of fulfilling the energy requirements associated with each EV $i \in \{1, \dots, V\}$, together with the power requirements associated with the AGC signal. While the nominal trajectory, calculated by DER-CAM, is a parameter for vehicle i at time k , a feasible trajectory is defined as any path for the SoC of a particular EV such that both physical and scheduling requirements are fulfilled. This means that for any vehicle i at time k , its SoC must lie within the interval $[\beta_i^-, \beta_i^+]$ to respect the limits of the battery. Likewise, each vehicle arrives at time a_i with a known SoC and is scheduled to leave at time d_i with a minimum level of energy for the EV to be able to operate normally, so

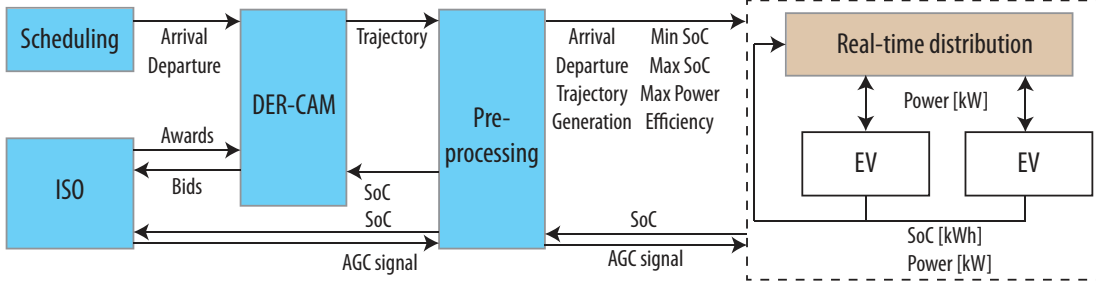


Fig. 1. Overview of the hierarchical control framework in the LAAFB project.

$E_{id_i} \in [E_i^-, E_i^+]$. For simplicity, $\beta_i^+ = E_i^+$ is assumed. There are power limits for the operation of each battery as well, so the charging/discharging (positive/negative) rate is $p_{ik} \in [m_i^-, m_i^+]$. This is illustrated in Figure 2.




Fig. 2. Battery model for EV i and feasible trajectories.

Definition 1: The laxity, ϕ_{ik} , is defined as the amount of time left until vehicle i must charge at its maximum charge rate to reach its minimum scheduled State of Charge (SoC) E_i^- , at departure time d_i .

$$\phi_{ik} = d_i - k - \frac{E_i^- - E_{ik}}{m_i^+} \quad (1)$$

It is common for battery models to consider an efficiency scalar $0 < \eta < 1$ to account for the difference between the power they received and the power they were able to transform into energy for storage. In the context of bidirectional charging, this effect must be considered both ways, as it is shown in Figure 3. When the power variable is positive, the batteries are




Fig. 3. Efficiency of the batteries when charging and discharging.

being charged and when it is negative, the batteries are being

discharged. Let x_{ik} be the power necessary from a source to charge the battery of vehicle i with p_{ik} , and p_{ik} be the power necessary from a battery to provide the grid with x_{ik} . This can be written in a compact way:

$$p_{ik} = \left(\frac{\eta_i^+}{2} + \frac{1}{2\eta_i^-} \right) x_{ik} + \left(\frac{\eta_i^+}{2} - \frac{1}{2\eta_i^-} \right) |x_{ik}| \quad (2)$$

The inverse relationship is defined $F(p_{ik}) = x_{ik}$ and can be easily derived from Equation 2.

Definition 2: A task T_i , can be represented by its parameters $(m_i^-, m_i^+, \eta_i^-, \eta_i^+, a_i, d_i, \beta^-, \beta^+, E_i^-, E_i^+)$, with states E_{ik} and ϕ_{ik} , as described in Table I. The index set of all active tasks in time step k is defined as $\mathbb{T}_k = \{i : k \in [a_i, d_i]\}$. Each active task models an EV that is available to provide regulation.

For notation simplicity, vectors are defined in bold symbols when referring to their components associated with active tasks: $\mathbf{x}_k = \{x_{ik}, i \in \mathbb{T}_k\}$, $\mathbf{p}_k = \{p_{ik}, i \in \mathbb{T}_k\}$, $\mathbf{E}_k = \{E_{ik}, i \in \mathbb{T}_k\}$, $\mathbf{r}_k = \{r_{ik}, i \in \mathbb{T}_k\}$, and $\Gamma_k^{+,-} = \{\Gamma_{ik}^{+,-}, i \in \mathbb{T}_k\}$.

B. Limits

If a task is close to its boundaries, depending on the energy state, it is possible that the charging rate may need to be reduced. The limits that guarantee that no upper or lower boundaries are violated for task i at time step k are denoted Γ_{ik}^+ and Γ_{ik}^- , respectively. These limits include all the information needed to ensure that the SoC stays within physical limits and is able to fulfill the task's minimum energy requirements [30].

$$\Gamma_{ik}^+ = \min \left[m_i^+, \frac{E_i^+ - E_{ik}}{\Delta t} \right] \quad (3)$$

$$\Gamma_{ik}^- = \min \left[\max \left[\max \left[m_i^-, \frac{\beta_i^- - E_{ik}}{\Delta t} \right], (1 - \phi_{ik}) m_i^+ \right], \Gamma_{ik}^+ \right] \quad (4)$$

The additional terms of Γ_{ik}^- , relative to Γ_{ik}^+ , are only relevant when $\phi_{ik} < 2$. When laxity is in that range, the lower bound of p_{ik} will be pushed in order to fulfill the minimum energy requirement at the time the task ceases to be active, and saturated by physical limits. Additionally, when aggregating these expressions for the fleet, the regulation capacity of the system can be calculated.

$$R_k^+ = \sum_{i \in \mathbb{T}_k} F(\Gamma_{ik}^+), \quad R_k^- = \sum_{i \in \mathbb{T}_k} F(\Gamma_{ik}^-) \quad (5)$$

Therefore, the *feasible regulation region* for time step k will be defined as the interval $[R_k^-, R_k^+]$.

C. Market Participation

LAAFB project is being developed using CAISO's rules for AS, which take into account both energy and AS when optimizing the system as a whole. Therefore, the EV fleet participates, in aggregate, in two markets: the energy market and the frequency regulation market². From these, the aggregate resource receives a frequency regulation signal every four seconds, which must be distributed among the individual vehicles. This AGC signal g_k for each time step k has two components: a fixed level of generation from the energy market award (g_k^{EM}), and a variable regulation quantity from the hourly AS market award (g_k^{FR}). The generation signal $g_k = g_k^{EM} + g_k^{FR}$ must be followed as accurately as possible for different magnitudes of the variable component.

The error is defined as the difference between the loads associated with the vehicles and the generation signal:

$$e_k = \sum_{i \in \mathbb{T}_k} x_{ik} - g_k \quad (6)$$

Thus, to maximize the performance, $|e_k|$ must be minimized. The relevant metric is the Accuracy, defined next.

$$\text{Accuracy} := 1 - \sum_{k=1}^T |e_k| / \sum_{k=1}^T |g_k^{FR}| \quad (7)$$

This metric is tied to payments for AS providers, which according to FERC's order 755 [36], are proportional to capacity and performance, which was already adopted by CAISO's AS market [37].

III. MYOPIC CONTROL

A myopic or short-sighted controller is presented, along with benchmarks that use simpler approaches. The ease of implementation is a relevant subject, so it will be useful to compare simulation results.

A. Trajectory Following (TF)

The proposed TF controller relies on previously calculated reference trajectories \mathbf{r}_k for the SoC of each EV in a fleet, considering their departures and arrivals, made by an external optimizer (DER-CAM [32]). It takes as an input the reference SoC trajectories, and reschedules at each operational time the actual charging trajectory for each vehicle in order to achieve the frequency regulation requirements, given by the realization of an AGC signal that tells the fleet which instantaneous aggregated power input/output it should have. These prespecified trajectories are calculated based on optimizing the participation of the fleet in the frequency regulation market and the charging of each EV under a retail electricity tariff. The

²While the resource settles its electricity costs at the retail price, in order to effectively participate in frequency regulation, the EV fleet must create a baseline electricity consumption on which it will regulate around, which is done in the wholesale energy market.

trajectories can be understood as a nominal path for the SoC of each vehicle, which enables the TF controller to incorporate information about future arrivals and departures.

For each time step k , the following convex optimization problem returns the optimal power vector, given the SoC of the previous time step \mathbf{E}_{k-1} . The proposed controller is defined as $\text{TF}(k) :=$

$$\min_{\mathbf{p}_k^c, \mathbf{p}_k^d} \alpha_1 \|\mathbf{r}_k - \mathbf{E}_k\|_2 + \alpha_2 |e_k| + \|\alpha_3 \circ \mathbf{p}_k^c\|_1 + \|\alpha_3 \circ \mathbf{p}_k^d\|_1$$

$$\text{s.t. } \mathbf{E}_k = \mathbf{E}_{k-1} + (\mathbf{p}_k^c + \mathbf{p}_k^d) \Delta t \quad (8)$$

$$e_k = \sum_{i \in \mathbb{T}_k} p_{ik}^c / \eta_i^+ + p_{ik}^d \eta_i^- - g_k \quad (9)$$

$$\Gamma_{k-1}^- \leq \mathbf{p}_k^c + \mathbf{p}_k^d \leq \Gamma_{k-1}^+ \quad (10)$$

$$\mathbf{p}_k^c \geq 0, \quad \mathbf{p}_k^d \leq 0 \quad (11)$$

Where \circ stands for element-wise vector multiplication, also known as the Hadamard product. In the formulation of TF controller, due to the non-convex relationship between \mathbf{p}_k and \mathbf{x}_k (Equation 2), we split the power variable in the EVs' side into charging (\mathbf{p}^c) and discharging (\mathbf{p}^d) power, given that one of them is 0. We relaxed the non-convex constraint $\mathbf{p}_k^c \cdot \mathbf{p}_k^d = 0$, and added a penalty for the variables so that the solution fulfills that requirement. This penalization also achieves non-aggressive control moves, thus reducing the cycling of the batteries compared to other benchmarks.

The real-time Trajectory Following (TF) controller consists of an objective function that sums three terms with different purposes: (1) tracking the SoC trajectories, (2) following the AGC signal and (3) penalizing the power variables for feasibility, with strictly positive penalties $\alpha_1, \alpha_2, \alpha_3$. As for the constraints, Eq. 8 represents the dynamics of the batteries, Eq. 9 represents the error in following the AGC signal and Eqs. 10 and 11 bound the decision variables.

If efficiencies η^+, η^- are 100%, the optimal solution always requires that $\mathbf{p}_k^c \cdot \mathbf{p}_k^d = 0 \quad \forall k$. When they are lower, it can be shown that the constraint also holds when a sufficient condition is fulfilled. Then, a convex relaxation based on tuning the penalties in the objective function is described in Theorem 1. Proof of this is presented in Appendix A.

Theorem 1. *If the penalties α_2, α_{3i} are such that:*

$$\alpha_{3i} > \alpha_2 \frac{1 - \eta_i^+ \eta_i^-}{2\eta_i^+} \quad \forall i,$$

the optimal solution to $\text{TF}(k)$ will satisfy $p_{ik}^c \cdot p_{ik}^d = 0 \quad \forall i \in \mathbb{T}_k$.

The rationale behind Theorem 1 is intuitive: there should be a threshold for the penalties of the control moves above which it is not optimal to charge and discharge simultaneously, because it would imply a higher cost for the objective function while not making the fleet response more accurate. If there was no penalty for \mathbf{p}^c and \mathbf{p}^d , the controller could do double-charging ($\mathbf{p}_k^c \cdot \mathbf{p}_k^d \neq 0$) if it implied better accuracy.

It should be noted that TF controller will provide a fast and feasible solution due to its convexity. Evidence suggests that if $\mathbf{p}_k^c \cdot \mathbf{p}_k^d = 0 \quad \forall k$ constraint is included by using binary variables instead of tuning the α_3 penalty according

to Theorem 1, the results of the non-convex problem are the same as in the convex problem, but computation time grows significantly.

A relevant subject for the implementation of the algorithm is how to tune the penalties. Due to the direct impact of $|e_k|$ on the Accuracy results, α_2 should be tuned as a large penalty relative to α_1 (given that E_i^- is always achieved).

B. Benchmarks

To benchmark the performance of the TF controller we consider two methods from the Processor Time Allocation (PTA) literature that have been applied to electric load scheduling [35], as well as version of the TF controller that assumes 100% round trip battery efficiency, all of which are applied in a sequential optimization simulation. In addition, a time-invariant benchmark was used.

1) *Earliest Deadline First (EDF)*: EDF creates a priority list based on the departure time of the tasks, and therefore will allow vehicles with the latest deadlines to remain at a low SoC until sufficient resources are available to charge them. We adapt this algorithm for discharging by coordinating vehicles such that those with the latest departure times are discharged first.

2) *Least Laxity First (LLF)*: LLF creates a merit order list sorted by laxity (see Eq. 1), and therefore will allow vehicles with larger laxity to remain at a low SoC until sufficient resources are available to charge them. Similarly to EDF, we adapt the algorithm for discharging such that the vehicles with the highest laxity are discharged first.

3) *Trajectory Following with approximate battery state (TFAPPROX)*: We remove the nonconvexity that results from bidirectional charging by making the approximation $p_{ik} \approx x_{ik}$ in the battery dynamics equation. However, the quality of the approximation degrades with declining efficiency.

4) *Time-invariant Trajectory Following (Oracle)*: We developed an *Oracle* benchmark that solves the complete run time at once. This additional benchmark provides a best-possible-performance case.

IV. MODEL PREDICTIVE CONTROL

This section will explain how to extend the TF myopic controller by implementing a predictive controller. We employ model predictive control, which takes into account not only the present current state of a system, but also its forecasted states over a finite time horizon (of length N), when making a decision. The underlying motivation is that MPC should allow the algorithm to achieve better Accuracy, because it will consider the EV arrivals and departures as well as a forecast of the AGC signal when deciding how to update the SoC of the vehicles.

A. Trajectory Following with Model Predictive Control (TFMPC)

A first approach to use MPC with TF would be to sum up the objective function values, while interpreting the bounds $\Gamma_{ik}^{+,-}$ as functions of the SoC. Including the upper bound

constraint with the future E_{ik} as a variable is not a problem, due to its concavity, but including the lower bound constraint would imply using a nonconvex expression as a lower bound:

$$p_{ik} \geq \min \left[\underbrace{\max \left[\max \left[m_i^-, \frac{\beta_i^- - E_{ik}}{\Delta t} \right], (1 - \phi_{ik}) m_i^+ \right]}_{\text{convex}}, \underbrace{\Gamma_{ik}^+}_{\text{concave}} \right] \quad (12)$$

In the myopic problem, this constraint's objective was to fulfill the task's minimum energy requirements at departure, which can also be achieved by transforming the laxity part of the lower bound for power into an energy constraint, so that the MPC problem is convex. For energy constraints to work, efficiency effects must be fully considered by the MPC controller in the time steps along the forecast horizon, so that the SoC can be properly estimated. Before formulating the MPC problem, some additional definitions are necessary:

- $\mathbb{F}_k = \{k, \dots, k + N\}$.
- $\hat{g}_{j|k}$ is the forecast for the AGC signal in period j , made in time step k .
- $\tilde{g}_j = \begin{cases} g_k, & \text{if } j = k \\ \hat{g}_{j|k}, & \text{if } j > k \end{cases}$
- Bold symbols must now include all tasks that may be active in the forecast horizon: $\{i \in \cup_{j \in \mathbb{F}_k} \mathbb{T}_j\}$.

For simulation purposes, the SoC of inactive tasks has to be updated using a previously calculated vector for the power variable. This represents the energy used by the EVs when they are not grid connected, and the values for E_i^- must be consistent with that vector. Thus, the MPC problem solved for each time step k , with feasible region \mathcal{Z}_k , is defined as $\text{TFMPC}(k) :=$

$$\min_{\mathbf{p}_j^c, \mathbf{p}_j^d} \sum_{j \in \mathbb{F}_k} \alpha_1 \|\mathbf{r}_j - \mathbf{E}_j\|_2 + \alpha_2 |e_j| + \|\alpha_3 \circ \mathbf{p}_j^c\|_1 + \|\alpha_3 \circ \mathbf{p}_j^d\|_1$$

$$\text{s.t. } \mathbf{E}_j = \mathbf{E}_{j-1} + (\mathbf{p}_j^c + \mathbf{p}_j^d) \Delta t \quad \forall j \in \mathbb{F}_k \quad (13)$$

$$e_j = \sum_{i \in \mathbb{T}_j} p_{ij}^c / \eta_i^+ + p_{ij}^d \eta_i^- - \tilde{g}_j \quad \forall j \in \mathbb{F}_k \quad (14)$$

$$\mathcal{Z}_k = \left\{ m_i^- \leq p_{ij}^c + p_{ij}^d \leq m_i^+, \right. \quad (15)$$

$$p_{ij}^c \geq 0, \quad p_{ij}^d \leq 0, \quad (16)$$

$$\beta_i^- \leq E_{ij} \leq \beta_i^+, \quad (17)$$

$$E_{ij} \geq E_i^- - (d_i - j) m_i^+ \Delta t \quad \forall i \in \mathbb{T}_j, \forall j \in \mathbb{F}_k \quad (18)$$

The terms in objective function of the real-time MPC controller have the same meaning as in TF. As shown in Theorem 1, properly tuning the penalties α_2, α_3 is critical for satisfying the $\mathbf{p}_k^c \cdot \mathbf{p}_k^d = 0$ constraint.

The constraints consider (1) the dynamics of the batteries (Eq. 13), (2) the present and future error in following the AGC signal (Eq. 14) and (3) the feasible set for both the power variables and the energy state (Eqs. 15 to 18). The purpose of Eq. 18 is twofold. First, it is needed to ensure that the task's minimum energy requirements are fulfilled. Second, it acts as a terminal constraint to ensure the feasibility of the MPC controller, regardless of the length of \mathbb{F}_k .

B. AGC signal forecast

As mentioned earlier, the generation signal g_k has a random component g_k^{FR} . As this paper handles uncertainty with the MPC approach, sequential forecasts are considered for g_k^{FR} by using an ARIMA approach with the information available up to each time step. This assumes that g_k^{FR} is a zero-mean signal, which may not be true in real applications (when defining drive cycles, non-zero means can be compensated [38]). To compare results, a comparison was made between simulations with an (1) an ARIMA forecast and a (2) a perfect forecast.

C. Operational capacity limits

There are some cases in which the input/output capacity of the fleet may be limited by exogenous system conditions. Let $\underline{C}_k, \bar{C}_k \in [0, 1]$ denote the lower and upper operational capacity limit at time step k , respectively, as a fraction of the original regulation capacity of the fleet. In order to account for this limitation in the MPC problem, the following constraint is included for each time step k :

$$C_k \sum_{i \in \{1, \dots, V\}} m_i^- \leq \sum_{i \in \mathbb{T}_k} p_{ik}^c + p_{ik}^d \leq \bar{C}_k \sum_{i \in \{1, \dots, V\}} m_i^+ \quad (19)$$

V. NUMERICAL RESULTS

The data set used in our numerical simulations is the same as in [30], with a fleet of 18 EVs, with a maximum regulation capacity of ± 15 kilowatts [kW] per vehicle, and a run time of two days with time steps every five minutes, during which the number of available vehicles changes according to a fixed task schedule. TFMPC used a forecast horizon of $N = 10$. Simulations were run using the MATLAB toolbox YALMIP [39] along with the solver Gurobi [40]. The variable component of the generation signal, g_k^{FR} , was simulated as an ARIMA time series for simulation purposes, based on historical data for PJM's *regd* test signal, meant to be used for fast regulation resources such as EV batteries [41]. For each time step, a forecast was made with the information of past realizations of the AGC signal. Input data were obtained from PJM's AS website [42], where normalized dynamic (*regd*) and traditional (*rega*) regulation signals are provided from seven days in May 2014. The dynamic regulation signal was used for this experiment. In terms of computation time, the simulations were run in on a 2.5 GHz Intel Core i5-3210M processor, and the TFMPC algorithm with $N = 10$ (the most computationally expensive) took always less than 0.1 seconds to be solved.

A. Accuracy results

All of the proposed versions of the algorithms were tested with six different test AGC signals and the Accuracy results were averaged. These represent the performance of each algorithm. Figure 4 shows the results for different magnitudes of the AGC signal and different battery efficiencies ($\eta^+ = \eta^- = \eta$). Note that the AGC Magnitude value is a higher bound for the signal's absolute value. Table II shows a comparison of the accuracy performance of each controller, relative to the Oracle case. The results must be discussed separately depending on the efficiency of the batteries. When

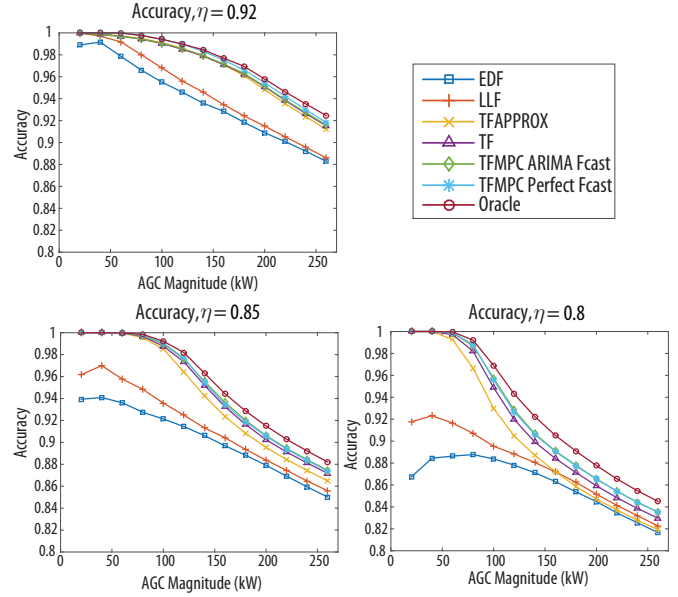


Fig. 4. Accuracy for different algorithms, AGC signal magnitudes and battery efficiencies $\eta \in \{0.92, 0.85, 0.8\}$. Mean for all seeds.

TABLE II
AVERAGE ACCURACY FOR EACH ALGORITHM, RELATIVE TO ORACLE

Oracle	TFMPC Perfect	TFMPC ARIMA	TF	TFAPPROX	LLF	EDF
100%	99.38%	99.31%	99.03%	98.43%	95.87%	94.61%

$\eta = 0.92$, the effective plugged in power capacity of each EV is ± 13.8 [kW], and sorting the performance of the algorithm and its benchmarks gives the following list: (1) Oracle, (2) TFMPC Perfect Forecast, (3) TFMPC ARIMA Forecast, (4) TF, (5) TFAPPROX, (6) LLF, and (7) EDF. The Accuracy results for LLF and EDF, the only algorithms that do not track predefined SoC trajectories, are noticeably worse than for the other algorithms, for all the magnitudes of the AGC signal; for the remainder of the algorithms the performance is similar.

When the efficiency of the batteries is decreased to $\eta = 0.85$, the effective plugged in power capacity of each EV is reduced to ± 12.75 [kW], and the performance of all the algorithms degrades but keeps the same order. In percentage terms, EDF and LLF are farther from the rest for low magnitudes of the AGC signal, and closer for high magnitudes. This happens because the benefits of trajectory following are greatest when real SoC trajectories are close to the reference trajectories – which happens to be when the AGC magnitude is smallest. On the other hand, when the AGC signal is large, the difference between real and reference trajectories is inevitably large – therefore reference trajectory following provides little benefit relative to EDF and LLF.

Finally, when $\eta = 0.8$, the effective plugged in power capacity of each EV is reduced to ± 12 [kW], and the tendency described in the former paragraph is confirmed: Earliest Deadline First (EDF) and Least Laxity First (LLF) show bad performance for low Automatic Generation Control (AGC) magnitudes, but their performance for high magnitudes compared to the other algorithms is similar, due to the limited

benefits to reference trajectory following when real trajectories are substantially different (as described in the previous paragraph). Note that both TFMPC options achieve Accuracy results that dominate over all the non-predictive algorithms. These are close to the Oracle's, but there is still some room to improve, which could be done with longer forecast horizons and more accurate forecasts.

1) *Effect of exogenous system limitations*: A case study was performed for the TFMPC ARIMA controller with the additional constraint (Eq. 19), where $C_k = \bar{C}_k = 0.5 \quad \forall k \in [20, 80]$, for $\eta = 0.92$. The results in Figure 5 show how the controller struggles to track the AGC signal during an episode of system limitations which translates into a tracking error. Mild congestion episodes for short periods of time could be

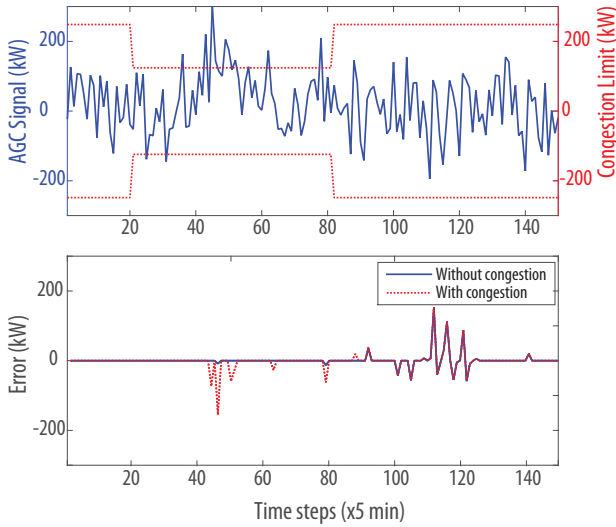



Fig. 5. Effect of exogenous system limitations. Limiting the capacity of the fleet to respond to the AGC signal reduces accuracy in the response.

an example of exogenous system limitations that the proposed AGC controllers can handle. However, in general, limitations associated with distribution congestion should be addressed by modifying the reference signal (g^{EM}), before the action of the AGC controller. In the case of the LAAF project, this task should be performed by the DER-CAM stage, as shown in Figure 1.

2) *Effect of delay*: Figure 6a shows the accuracy of the different control algorithms for the case with a delay of 1/10 of an AGC time-step. It can be observed that the accuracy is reduced between 10-15% for all the algorithms, and for different magnitudes of the AGC signal, when compared with the case without delay (see Figure 4). Nevertheless, the results still show a superior performance of MPC-based controllers over non-predictive algorithms for the case of delayed fleet response. An improved performance can be obtained using the predictive feature of MPC controllers to anticipate the AGC signal for a known, constant delay. In particular, Figure 6b shows the case in which the TFMPC ARIMA controller anticipates the delay of the response to the AGC signal by implementing its predicted response, achieving approximately a 3% increase in accuracy with respect to the non-anticipative case.



(a) Accuracy of different algorithms with delay (b) Accuracy of anticipative TFMPC ARIMA with delay

Fig. 6. Delay effect on the response of the fleet. (a) Results for different algorithms, AGC signal magnitudes and fixed battery efficiency $\eta = 0.92$. Mean for all seeds. (b) Accuracy achieved by MPC fast ARIMA algorithm, for different AGC signal magnitudes and fixed battery efficiency $\eta = 0.92$. Mean for all seeds.

B. Regulation capacity results

We define the feasible regulation region $[R_k^-, R_k^+]$ as the range of power in which a generation signal g_k should lie so that the fleet can provide regulation with no error, while being able to fulfill the minimum energy requirements.

Outside the feasible region, the algorithms with a myopic approach behave differently than TFMPC. For myopic algorithms, their behavior is easy to understand: in that situation, all the difference between g_k and the fleet's capacity to provide regulation results in error. In contrast, the look-ahead characteristic of the latter provides other possibility. TFMPC algorithms may choose to save battery energy in a given time step, resulting in avoidable error in the short run in favor of reducing long-run error to minimize total error. This emphasizes the relevance of properly tuning the parameters $\alpha_1, \alpha_2, \alpha_3$ in a way that the system makes desirable decisions in the face of these trade-offs. The results for R^+ and R^- are shown in Figure 7.

When sorting the algorithms by the width of the mean feasible regulation region they achieved, in general the order is the same as in the ranking shown for Accuracy (not including Oracle).

As for R^+ results (charging capacity), TF approaches achieve better levels than EDF and LFF in all cases, which is another benefit of tracking SoC trajectories. As the system receives AGC signals with higher magnitude or uses less-efficient batteries, batteries get drained, and therefore R^+ increases because there is more room for the batteries to get charged.

On the other hand, R^- (discharging capacity) is the real bottleneck of the system for large AGC signals, because the battery SoC is typically well below what is required to follow the AGC signals. Furthermore, minimum energy requirements at a task's departure also constrain the discharging capacity of each EV. Results for R^- are similar to the Accuracy results when sorting the performance of the algorithms, except for TFAPPROX when $\eta = 0.8$; the bad quality of the $p_{ik} \approx x_{ik}$ approximation directly impacts the discharging capacity of the batteries. It is clear that implementing MPC improves the capacity of the EV fleet to discharge its batteries, and this effect is magnified and therefore can be seen more clearly

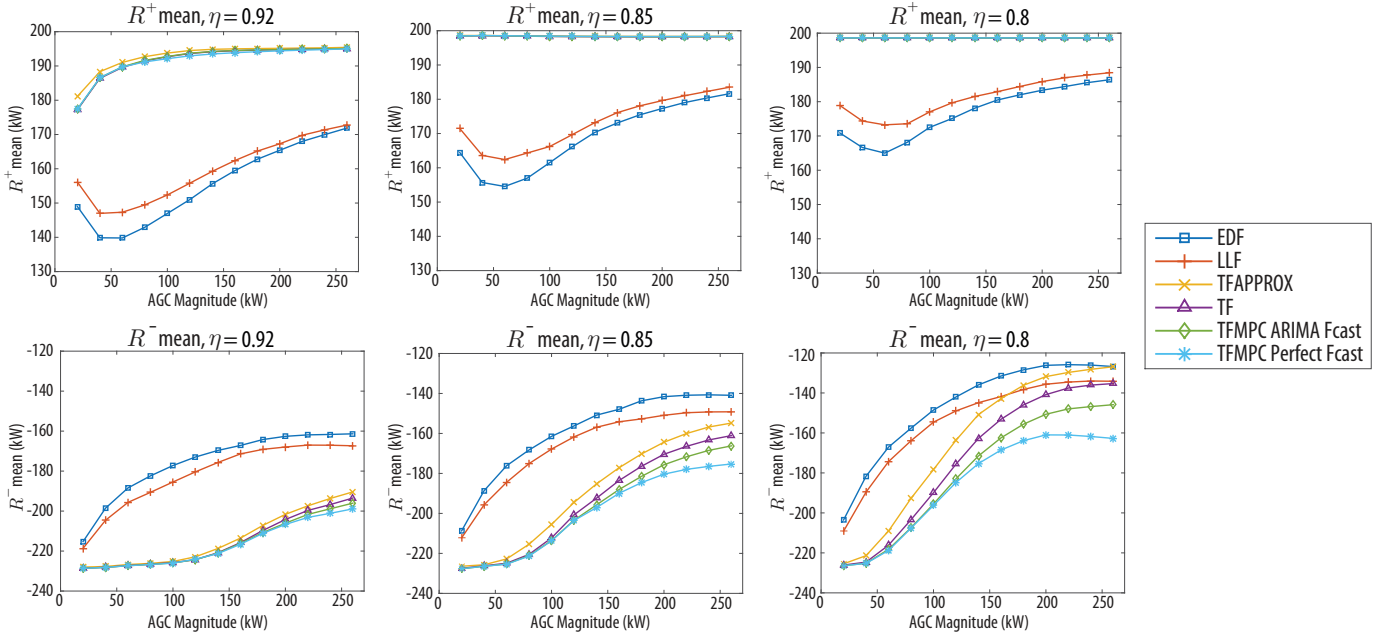


Fig. 7. Mean R^+ and R^- for different algorithms, AGC signal magnitudes and battery efficiencies $\eta \in \{0.92, 0.85, 0.8\}$.

when the efficiency of the batteries decreases. Thus, the importance of using MPC as opposed to myopic strategies is greater when the system works closer to its limits.

C. Cycling results

Here we examine the impact of the control approaches on cycling – which is believed to degrade battery state of health – by using the arc length of the SoC curves as a proxy. Table III shows the average value of this metric over all the experiments, relative to the results of EDF. TF approaches

TABLE III

AVERAGE ARC LENGTH FOR EACH ALGORITHM, RELATIVE TO EDF

TFMPC Perfect	TFMPC ARIMA	TF	TFAPPROX	LLF	EDF
90.63%	91.04%	92.24%	96.07%	99.78%	100%

lead to less cycling, and we attribute this result to one key factor: by penalizing deviations from an SoC trajectory, each of the TF approaches tend to cycle all batteries in roughly the same way. This results in roughly equal distribution of ramping across all the EVs, rather than distributing power changes to EVs in the most extreme states (as with EDF and LLF). We also see that, as one might expect, appropriately penalizing the control variables (as with TF vs TFAPPROX) results in better performance.

D. Individual AGC signal and EV results

Lastly, we show some results for seed 5, when $\eta^+ = \eta^- = 0.85$ and the AGC signal magnitude is of 100 kW. Figure 8 shows the varying EV availability, the AGC signal g_k with its two components, and the error in the EV fleet response e_k . Noteworthy, the error's magnitude is higher when fewer EVs are available due to a reduced regulation capacity. However,

this error is reduced when more complex algorithms based on convex optimization and MPC are used to track the AGC signal. Figure 9 shows how the SoC of a particular EV changes

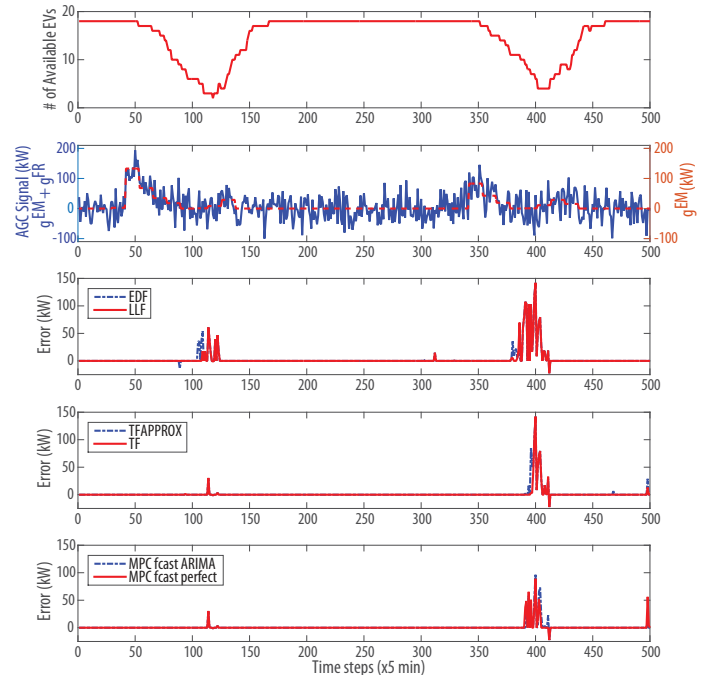


Fig. 8. EV availability, AGC Signal and response error for different strategies.

during the simulation with each of the strategies. In this particular case, the arc length of the SoC curves is noticeably longer for EDF and LLF than for the rest of the strategies. It can also be noted that the SoC trajectories were not followed perfectly, given that the higher weight was assigned to α_2 in order to prioritize following the complete AGC signal (and therefore achieve better accuracy). Nevertheless, the minimum

energy requirements are hard constraints in the trajectory-following controllers; thus, they are always achieved.




Fig. 9. State of Charge of EV 13 for different strategies.

VI. CONCLUSIONS

In this paper, we focus on the design of charging schedules of EVs for the provision of frequency regulation services. In particular, we propose several real-time scheduling schemes differentiated by the way of handling future information (myopic/non-myopic), the level of accuracy on following regulation signals and the resulting cycling on the batteries. A method for considering bidirectional efficiency while enabling the estimation of the future state of charge of the EVs in the fleet is provided, which allows the use of model predictive control schemes.

Extensive simulation results show the trade-off between the complexity of the controllers and their accuracy on following regulation signals: for practical implementations, both the ease of use and the performance are relevant. A key insight is that higher accuracy in following regulation signals coincides with less cycling of the batteries and, in most cases, with better regulation capacity. This highlights the importance of keeping the state of charge of the batteries away from their physical boundaries when providing frequency regulation services.

The generality of the approach enables the use of the same framework for any kind of energy battery, such as water reservoirs or HVAC loads. Other stochastic fast-response resources suitable for demand response, such as buildings or industrial processes, can be integrated into the proposed controller as well. Future work should take into account uncertainty in the EVs' arrival and departure, in which MPC along with advanced forecasting techniques can be specially valuable. Moreover, a control formulation explicitly considering delays in the EV fleet response should also be taken into account.

ACKNOWLEDGEMENTS

The authors would like to thank Doug Black, PI of the LAAFB project, and Somil Bansal for insightful discussions, and CONICYT for partially funding George Wenzel's research stay at UC Berkeley.

APPENDIX

A. Proofs

Proof: Theorem 1. Assume that the optimal solution for some i, k is given by $p_{ik}^c = t_1 + t_2, p_{ik}^d = -t_2$ such that $t_1 > 0, t_2 > 0$ (so that $(t_1 + t_2)t_2 \neq 0$). For simplicity, also assume that $\eta_i^+ = \eta^+, \eta_i^- = \eta^- \forall i$, and let $f(\cdot)$ denote the objective function. Condition $p_{ik}^c \cdot p_{ik}^d = 0$ would always hold if a sufficient condition was found so that $f(t_1, 0) < f(t_1 + t_2, -t_2)$. When comparing $f(t_1 + t_2, -t_2)$ with $f(t_1, 0)$, the first term, $\alpha_1 \|\mathbf{r}_k - \mathbf{E}_k\|_2$, is the same in both cases. Thus, the rest of $f(\cdot)$ remains to be compared. When replacing, the following expression is obtained:

$$\begin{aligned} f(t_1, 0) &< f(t_1 + t_2, -t_2) \\ \alpha_2 \left| \frac{t_1}{\eta^+} - g_k \right| + \alpha_{3i} |t_1| &< \alpha_2 \left| \frac{t_1 + t_2}{\eta^+} - t_2 \eta^- - g_k \right| + \alpha_{3i} (|t_1 + t_2| + |-t_2|) \\ \Leftrightarrow \alpha_2 \left| \frac{t_1}{\eta^+} - g_k \right| &< \alpha_2 \left| \frac{t_1}{\eta^+} + t_2 \left(\frac{1}{\eta^+} - \eta^- \right) - g_k \right| + 2\alpha_{3i} t_2 \end{aligned} \quad (\text{A.1})$$

Note that because of the assumption on t_2 and $0 < \eta < 1$, $t_2 \left(\frac{1}{\eta^+} - \eta^- \right)$ is a strictly positive term. Some cases must be analyzed:

$$1) \frac{t_1}{\eta^+} - g_k \geq 0:$$

$$\begin{aligned} \alpha_2 \left(\frac{t_1}{\eta^+} - g_k \right) &< \alpha_2 \left(\frac{t_1}{\eta^+} + t_2 \left(\frac{1}{\eta^+} - \eta^- \right) - g_k \right) + 2\alpha_{3i} t_2 \\ \Leftrightarrow 0 &< \alpha_2 t_2 \left(\frac{1}{\eta^+} - \eta^- \right) + 2\alpha_{3i} t_2 \end{aligned} \quad (\text{A.2})$$

Which holds because all the terms are positive.

$$2) \frac{t_1}{\eta^+} - g_k < 0:$$

This case must be split into three more cases:

$$a) \frac{t_1}{\eta^+} + t_2 \left(\frac{1}{\eta^+} - \eta^- \right) - g_k > 0$$

$$\begin{aligned} -\alpha_2 \left(\frac{t_1}{\eta^+} - g_k \right) &< \alpha_2 \left(\frac{t_1}{\eta^+} + t_2 \left(\frac{1}{\eta^+} - \eta^- \right) - g_k \right) + 2\alpha_{3i} t_2 \\ \Leftrightarrow 0 &< 2\alpha_2 \left(\frac{t_1}{\eta^+} + t_2 \left(\frac{1}{\eta^+} - \eta^- \right) - g_k \right) \\ &+ 2\alpha_{3i} t_2 - \alpha_2 t_2 \left(\frac{1}{\eta^+} - \eta^- \right) \end{aligned} \quad (\text{A.3})$$

Because of the assumption,

$$2\alpha_2 \left(\frac{t_1}{\eta^+} + t_2 \left(\frac{1}{\eta^+} - \eta^- \right) - g_k \right) > 0 \quad (\text{A.4})$$

So, a condition can be obtained if the rest is also positive:

$$\begin{aligned} 2\alpha_{3i} t_2 - \alpha_2 t_2 \left(\frac{1}{\eta^+} - \eta^- \right) &> 0 \\ \Leftrightarrow \alpha_{3i} &> \alpha_2 \frac{1 - \eta^+ \eta^-}{2\eta^+} \end{aligned} \quad (\text{A.5})$$

$$b) \frac{t_1}{\eta^+} + t_2 \left(\frac{1}{\eta^+} - \eta^- \right) - g_k = 0$$

$$-\alpha_2 \left(\frac{t_1}{\eta^+} - g_k \right) < 2\alpha_{3i} t_2 \quad (\text{A.6})$$

Because of the assumption,

$$\alpha_2 t_2 \left(\frac{1}{\eta^+} - \eta^- \right) < 2\alpha_{3i} t_2 \quad (\text{A.7})$$

$$\Leftrightarrow \alpha_{3i} > \alpha_2 \frac{1 - \eta^+ \eta^-}{2\eta^+}$$

$$c) \frac{t_1}{\eta^+} + t_2 \left(\frac{1}{\eta^+} - \eta^- \right) - g_k < 0$$

$$\begin{aligned} -\alpha_2 \left(\frac{t_1}{\eta^+} - g_k \right) &< -\alpha_2 \left(\frac{t_1}{\eta^+} + t_2 \left(\frac{1}{\eta^+} - \eta^- \right) - g_k \right) + 2\alpha_{3i} t_2 \\ \Leftrightarrow 0 &< -\alpha_2 t_2 \left(\frac{1}{\eta^+} - \eta^- \right) + 2\alpha_{3i} t_2 \\ \Leftrightarrow \alpha_{3i} &> \alpha_2 \frac{1 - \eta^+ \eta^-}{2\eta^+} \end{aligned} \quad (\text{A.8})$$

Therefore, if the sufficient condition is respected for every i , $p_{ik}^c \cdot p_{ik}^d = 0$ will be satisfied for every i, k in the optimal solution. ■

REFERENCES

- [1] N. G. Gast, J.-Y. Le Boudec, A. Proutiere, and D.-C. Tomozei, "Impact of Storage on the Efficiency and Prices in Real-Time Electricity Markets," EPFL, Switzerland, Tech. Rep., 2013.
- [2] D. L. Anderson, "An evaluation of current and future costs for lithium-ion batteries for use in electrified vehicle powertrains," Ph.D. dissertation, Duke University, 2009.
- [3] B. Nykvist and M. Nilsson, "Rapidly falling costs of battery packs for electric vehicles," *Nature Climate Change*, vol. 5, no. 4, pp. 329–332, 2015.
- [4] M. Alizadeh, M. P. Moghaddam, N. Amjadi, P. Siano, and M. Sheikh-El-Eslami, "Flexibility in future power systems with high renewable penetration: A review," *Renewable and Sustainable Energy Reviews*, vol. 57, pp. 1186 – 1193, 2016.
- [5] S. Meyn, M. Negrete-Pincetic, G. Wang, A. Kowli, and E. Shafiepoorfar, "The value of volatile resources in electricity markets," in *Proceedings of the 49th IEEE Conference on Decision and Control (CDC)*, Dec 2010, pp. 1029–1036.
- [6] A. Papavasiliou, S. Oren, and R. O'Neill, "Reserve requirements for wind power integration: A scenario-based stochastic programming framework," *Power Systems, IEEE Transactions on*, vol. 26, no. 4, pp. 2197–2206, 2011.
- [7] J. Cochran, M. Miller, O. Zinaman, M. Milligan, D. Arent, B. Palmintier, M. O'Malley, S. Mueller, E. Lannoye, A. Tuohy, B. Kujala, M. Sommer, H. Holttinen, J. Kiviluoma, and S. Soonee. (2014) Flexibility in 21st century power systems. [Online]. Available: <http://www.nrel.gov/docs/fy14osti/61721.pdf>
- [8] A. Nayyar, M. Negrete-Pincetic, K. Poolla, and P. Varaiya, "Duration-differentiated energy services with a continuum of loads," *IEEE Transactions on Control of Network Systems*, vol. 3, no. 2, pp. 182–191, June 2016.
- [9] 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.
- [10] W. Kempton, V. Udo, K. Huber, K. Komara, S. Letendre, S. Baker, D. Brunner, and N. Pearre. (2008) A test of vehicle-to-grid (v2g) for energy storage and frequency regulation in the pjim system. [Online]. Available: <http://www1.udel.edu/V2G/resources/test-v2g-in-pjim-jan09.pdf>
- [11] PowerUp. (2013) Power up project. [Online]. Available: <http://www.power-up.org/>
- [12] K. Mets, T. Verschueren, W. Haerick, C. Devellder, and F. De Turck, "Optimizing smart energy control strategies for plug-in hybrid electric vehicle charging," in *Proceedings of the Network Operations and Management Symposium Workshops (NOMS Wksp), 2010 IEEE/IFIP*, April 2010, pp. 293–299.
- [13] J. Tomić and W. Kempton, "Using fleets of electric-drive vehicles for grid support," *Journal of Power Sources*, vol. 168, no. 2, pp. 459 – 468, 2007.
- [14] L. Göransson, S. Karlsson, and F. Johnsson, "Integration of plug-in hybrid electric vehicles in a regional wind-thermal power system," *Energy Policy*, vol. 38, no. 10, pp. 5482 – 5492, 2010.
- [15] J. Pillai and B. Bak-Jensen, "Electric vehicle based battery storages for future power system regulation services," in *Proceedings of the 5th Nordic Wind Power Conference*. Technical University of Denmark (DTU), 2009.
- [16] C. T. Li, C. Ahn, H. Peng, and J. Sun, "Synergistic control of plug-in vehicle charging and wind power scheduling," *IEEE Transactions on Power Systems*, vol. 28, no. 2, pp. 1113–1121, May 2013.
- [17] J. Lin, K. C. Leung, and V. O. K. Li, "Optimal scheduling with vehicle-to-grid regulation service," *IEEE Internet of Things Journal*, vol. 1, no. 6, pp. 556–569, Dec 2014.
- [18] J. Donadee and M. D. Ili, "Stochastic optimization of grid to vehicle frequency regulation capacity bids," *IEEE Transactions on Smart Grid*, vol. 5, no. 2, pp. 1061–1069, March 2014.
- [19] H. Liu, Z. Hu, Y. Song, and J. Lin, "Decentralized vehicle-to-grid control for primary frequency regulation considering charging demands," *IEEE Transactions on Power Systems*, vol. 28, no. 3, pp. 3480–3489, Aug 2013.
- [20] 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.
- [21] C. Wu, H. Mohsenian-Rad, and J. Huang, "Vehicle-to-aggregator interaction game," *IEEE Transactions on Smart Grid*, vol. 3, no. 1, pp. 434–442, March 2012.
- [22] J. J. Escudero-Garzas, A. Garcia-Armada, and G. Seco-Granados, "Fair design of plug-in electric vehicles aggregator for v2g regulation," *IEEE Transactions on Vehicular Technology*, vol. 61, no. 8, pp. 3406–3419, Oct 2012.
- [23] E. Yao, V. W. S. Wong, and R. Schober, "Robust frequency regulation capacity scheduling algorithm for electric vehicles," *IEEE Transactions on Smart Grid*, vol. PP, no. 99, pp. 1–14, 2016.
- [24] H. Liu, Z. Hu, Y. Song, J. Wang, and X. Xie, "Vehicle-to-grid control for supplementary frequency regulation considering charging demands," *IEEE Transactions on Power Systems*, vol. 30, no. 6, pp. 3110–3119, Nov 2015.
- [25] S. I. Vagropoulos, D. K. Kyriazidis, and A. G. Bakirtzis, "Real-time charging management framework for electric vehicle aggregators in a market environment," *IEEE Transactions on Smart Grid*, vol. 7, no. 2, pp. 948–957, March 2016.
- [26] C. L. Floch, E. Kara, and S. Moura, "Pde modeling and control of electric vehicle fleets for ancillary services: A discrete charging case," *IEEE Transactions on Smart Grid*, vol. PP, no. 99, pp. 1–1, 2016.
- [27] S. Sun, M. Dong, and B. Liang, "Real-time welfare-maximizing regulation allocation in dynamic aggregator-avs system," *IEEE Transactions on Smart Grid*, vol. 5, no. 3, pp. 1397–1409, May 2014.
- [28] R. J. Bessa and M. A. Matos, "Economic and technical management of an aggregation agent for electric vehicles: a literature survey," *European Transactions on Electrical Power*, vol. 22, no. 3, pp. 334–350, 2012. [Online]. Available: <http://dx.doi.org/10.1002/etep.565>
- [29] B. Xu, Y. Dvorkin, D. S. Kirschen, C. Silva-Monroy, and J.-P. Watson, "A comparison of policies on the participation of storage in us frequency regulation markets," *arXiv preprint arXiv:1602.04420*, 2016.
- [30] F. Juul, M. Negrete-Pincetic, J. MacDonald, and D. Callaway, "Real-time scheduling of electric vehicles for ancillary services," in *Proceedings of the 2015 IEEE Power Energy Society General Meeting*, July 2015, pp. 1–5.
- [31] M. Zeilinger, "Real-time Model Predictive Control," Ph.D. dissertation, ETH Zurich, ETH Zurich, Sep. 2011. [Online]. Available: <https://control.ee.ethz.ch/index.cgi?page=publications;action=details;id=3977>
- [32] C. Marnay, T. W. Chan, N. DeForest, J. Lai, J. MacDonald, M. Stadler, T. Erdmann, A. Hoheisel, M. Mueller, S. Sabre, E. Koch, P. Lipkin, R. W. Anderson, S. Gerber, and E. Reid, "Los angeles air force base vehicle to grid pilot project." ECEEE, 2013.
- [33] Lawrence Berkeley National Lab. DER-CAM. [Online]. Available: <https://building-microgrid.lbl.gov/projects/der-cam>
- [34] Kisenum. (2016) Connecting electric vehicles and storage to the smart grid. [Online]. Available: <http://www.kisenum.com/#home>
- [35] A. Subramanian, M. Garcia, A. Dominguez-García, D. Callaway, K. Poolla, and P. Varaiya, "Real-time scheduling of deferrable electric loads," in *Proceedings of the 2012 American Control Conference (ACC)*, June 2012, pp. 3643–3650.
- [36] FERC. (2011) Order 755: Frequency Regulation Compensation in the Organized Wholesale Power Markets. [Online]. Available: <http://www.ferc.gov/whats-new/comm-meet/2011/102011/E-28.pdf>
- [37] CAISO. (2015) Business practice manual for market operations, version 45. [Online]. Available: https://bpmcm.caiso.com/BPM%20Document%20Library/Market%20Operations/Market%20Operations%20BPM%20V45_redline.pdf
- [38] R. P. Hafen, V. V. Vishwanathan, K. Subbarao, and M. C. Kintner-Meyer, *Requirements for Defining Utility Drive Cycles: An Exploratory Analysis of Grid Frequency Regulation Data for Establishing Battery Performance Testing Standards*. Pacific Northwest National Laboratory, Oct 2011. [Online]. Available: <http://www.osti.gov/scitech/servlets/purl/1028571>
- [39] J. Löfberg, "Yalmip : A toolbox for modeling and optimization in MATLAB," in *Proceedings of the CACSD Conference*, Taipei, Taiwan, 2004. [Online]. Available: <http://users.isy.liu.se/johanl/yalmip>
- [40] Gurobi Optimization, Inc., "Gurobi optimizer reference manual," 2015. [Online]. Available: <http://www.gurobi.com>
- [41] PJM. (2013) PJM ancillary services - regulation. [Online]. Available: https://pjm.adobeconnect.com/_a16103949/p7s5g501jfn/
- [42] ——. (2015) PJM ancillary services. [Online]. Available: <http://www.pjm.com/markets-and-operations/ancillary-services.aspx>



George Wenzel was born in Santiago, Chile. He received a B.Sc. and M.Sc. in Electrical Engineering from Pontificia Universidad Católica de Chile, in 2014 and 2016, respectively. He is currently a R&D engineer at the Energy Optimization, Control and Markets Laboratory (OCM-Lab) at Pontificia Universidad Católica de Chile. His research activities include electricity markets, demand response architecture and vehicle-to-grid control.



Duncan S. Callaway (M'08) received the B.S. degree in mechanical engineering from the University of Rochester, Rochester, NY, in 1995, and the Ph.D. degree in theoretical and applied mechanics from Cornell University, Ithaca, NY, in 2001. He is currently an Assistant Professor of Energy and Resources and Mechanical Engineering, University of California, Berkeley. Prior to joining the University of California, he was first an NSF Postdoctoral Fellow with the Department of Environmental Science and Policy, University of California, Davis, subsequently worked as a Senior Engineer at Davis Energy Group, Davis, CA, and PowerLight Corporation, Berkeley CA, and was most recently a Research Scientist with the University of Michigan. His current research interests include the areas of power management, modeling and control of aggregated storage devices, spatially distributed energy resources, and environmental impact assessment of energy technologies.



Matías Negrete-Pincetic received a B.Sc. in Electrical Engineering and a M.Sc. in Physics from Pontificia Universidad Católica de Chile, and a M.Sc. in Physics and a Ph.D. in Electrical and Computer Engineering from the University of Illinois at Urbana-Champaign, USA. He was a Postdoctoral Associate at the University of California, Berkeley, CA, USA. Currently, he is an Assistant Professor at the Electrical Engineering Department at Pontificia Universidad Católica de Chile. His current research activities include operation, control and planning of energy systems, stochastic control, electricity market design and energy policy.



Daniel E. Olivares (S'11-M'14) was born in Santiago, Chile. He received the B.Sc. and Engineer degrees in electrical engineering from the University of Chile, Santiago, in 2006 and 2008, respectively, and the Ph.D. degree in electrical and computer engineering from the University of Waterloo, Waterloo, ON, Canada, in 2014.

He is currently an Assistant Professor with the Department of Electrical Engineering, Pontificia Universidad Católica de Chile, Santiago. His current research interests include modeling, simulation, and control and optimization of power systems in the context of smart grids.



Jason MacDonald (S'14) is a Scientific Research Associate at Lawrence Berkeley National Laboratory, and a Ph.D. student in Energy and Resources at the University of California, Berkeley. He received an MS in Natural Resources and the Environment and an MSE in Mechanical Engineering at the University of Michigan, Ann Arbor in 2011. His research is centered on electricity market participation of distributed energy resources, with particular emphasis in the nexus of market interaction and controls. His work history includes engineering and

research in plug-in electric vehicles, photovoltaic systems design, and demand response for ancillary services.